ISCR

Institute for Scientific Computing Research

Annual Report: Fiscal Year 2004

Collaborating with Academia

Lawrence Livermore National Laboratory
P.O. Box 808, L-561, Livermore, CA 94551

The University Relations Program (URP) encourages collaborative research between Lawrence Livermore National Laboratory (LLNL) and the University of California campuses. The Institute for Scientific Computing Research (ISCR) actively participates in such collaborative research, and this report details the Fiscal Year 2004 projects jointly served by URP and ISCR. For a full discussion of all URP projects in FY 2004, please request a copy of the URP FY 2004 Annual Report by contacting

Lawrence Livermore National Laboratory
Edie Rock, University Relations Program
P. O. Box 808, L-413
Livermore, CA 94551

UCRL-TR-210349

Work performed under t -7405-ENG-48.

Contents

Director's Report

Large-scale scientific computation and all of the disciplines that support and help to validate it have been placed at the focus of Lawrence Livermore National Laboratory (LLNL) by the Advanced Simulation and Computing (ASC) program of the National Nuclear Security Administration (NNSA) and the Scientific Discovery through Advanced Computing (SciDAC) initiative of the Office of Science of the Department of Energy (DOE). The maturation of computational simulation as a tool of scientific and engineering research is underscored in the November 2004 statement of the Secretary of Energy that, "high performance computing is the backbone of the nation's science and technology enterprise."

LLNL operates several of the world's most powerful computers—including today's single most powerful—and has undertaken some of the largest and most compute-intensive simulations ever performed. Ultrascale simulation has been identified as one of the highest priorities in DOE's facilities planning for the next two decades. However, computers at architectural extremes are notoriously difficult to use efficiently. Furthermore, each successful terascale simulation only points out the need for much better ways of interacting with the resulting avalanche of data. Advances in scientific computing research have, therefore, never been more vital to LLNL's core missions than at present. Computational science is evolving so rapidly along every one of its research fronts that to remain on the leading edge, LLNL must engage researchers at many academic centers of excellence. In Fiscal Year 2004, the Institute for Scientific Computing Research (ISCR) served as one of LLNL's main bridges to the academic community with a program of collaborative subcontracts, visiting faculty, student internships, workshops, and an active seminar series.

The ISCR identifies researchers from the academic community for computer science and computational science collaborations with LLNL and hosts them for short- and long-term visits with the aim of encouraging long-term academic research agendas that address LLNL's research priorities. Through such collaborations, ideas and software flow in both directions, and LLNL cultivates its future workforce. The Institute strives to be LLNL's "eyes and ears" in the computer and information sciences, keeping the Laboratory aware of and connected to important external advances. It also attempts to be the "feet and hands" that carry those advances into the Laboratory and incorporates them into practice.

ISCR research participants are integrated into LLNL's Computing and Applied Research (CAR)

Department, especially into its Center for Applied Scientific Computing (CASC). In turn, these organizations address computational challenges arising throughout the rest of the Laboratory. Administratively, the ISCR flourishes under LLNL's University Relations Program (URP). Together with the other five institutes of the URP, it navigates a course that allows LLNL to benefit from academic exchanges while preserving national security. While it is difficult to operate an academic-like research enterprise within the context of a national security laboratory, the results declare the challenges well met and worth the continued effort.

Fiscal year 2004 was the fifth full year under Acting Director David Keyes. Keyes, the Fu Foundation Professor of Applied Mathematics at Columbia University and an ISCR faculty participant since October 1997, dedicated one-third of his time to the technical program of the ISCR. James McGraw continued as the Deputy Director of the ISCR, and Linda Becker as the Institute Administrator. Paula Ashley, Pamela Mears, and Char Paulo logistically supported the large visitor and summer programs of the ISCR.

The ISCR continues to have a small contingent of research staff members within its organization. Three ISCR staff—Nelson Max, Garry Rodrigue, and Rao Vemuri—hold joint appointments as professors at the University of California, Davis and senior researchers at LLNL. In addition, the ISCR hosted eight post-doctoral staff: Alison Baker, David Buttler, Tzanio Kolev, Shawn Newsam, Stefan Nilsson, Dan Reynolds, Megan Thomas, and Qing Yi. Finally, the ISCR served as the host for 10 students (listed in Table 1) on a Student-Employee Graduate Fellowship (SEGRF). This fellowship enables students to work with LLNL researchers half-time while pursuing their PhDs.

The ISCR enables substantial interactions between academia and LLNL staff through consultants and participating guests. Consulting agreements are vehicles for permitting academics to interact with LLNL in a compensated fashion. Consultants may serve on review committees, present short courses, and visit LLNL periodically for

Table 1. Students on a Student-Employee Graduate Fellowship (SEGRF).

Name	University	LLNL Advisor(s)	Time at LLNL
Peer-Timo Bremer	University of California, Davis	Dan Laney	June 12, 2002 – Aug. 6, 2004
Sam Brockington	University of California, Davis	Garry Rodrigue & Dave Hwang	Oct. 8, 2001 – Aug. 31, 2006
Paul Castellucci	Stanford University	Rose McCallen	Oct. 13, 2003 – Jan. 31, 2005
Aaron Fisher	University of California, Davis	Garry Rodrigue	July 1, 2002 – June 30, 2006
Benjamin Gregorski	University of California, Davis	Mark Duchaineau	June 25, 2001 – Aug. 13, 2004
Jeff Hagelberg	Purdue University	Paul Amala	Sept. 22, 2003 – Sept. 25, 2004
Aaron Herrnstein	University of California, Davis	Michael Wickett	Mar. 26, 2004 – Mar. 31, 2004
Andrew Nonaka	University of California, Davis	David Trebotich	Oct. 1, 2003 – Sept. 30, 2007
Rob Rieben	University of California, Davis	Garry Rodrigue	Oct. 2, 2000 – Oct. 1, 2004
Joshua Senecal	University of California, Davis	Mark Duchaineau	Nov. 1, 2001 – Oct. 31, 2005

Table 2. Nine ISCR consultants for FY 2004.

Consultant	Affiliation	LLNL Contact
Randolph Bank	University of California, San Diego	Rob Falgout
Achiezer Brandt	University of California, Los Angeles	Rob Falgout
Gene Golub	Stanford University	Edmond Chow
Anne Greenbaum	University of Washington	Peter Brown
Heinz-Otto Kreiss	University of California, Los Angeles	Lori Diachin & Bill Henshaw
Thomas Manteuffel	University of Colorado	Rob Falgout & Peter Brown
Stephen McCormick	University of Colorado	Rob Falgout & Peter Brown
Linda Petzold	University of California, Santa Barbara	Carol Woodward & Radu Serban
Homer Walker	Worcester Polytech Institute	Peter Brown

technical meetings. All consultants have a specific LLNL technical point of contact for overseeing their interactions. Table 2 lists the nine ISCR consultants for FY 2004.

Participating Guests are researchers from academia or industry that need intermittent access to LLNL staff and facilities. This status permits an appropriate security clearance and the ability to quickly arrange for on-site visits with LLNL staff over a period of one month to two years. Table 3 lists ISCR's 38 participating guests for FY 2004.

The pages of this annual report summarize the activities of the faculty members, postdoctoral researchers, students, and guests from industry and other laboratories who participated in LLNL's computational mission under the auspices of the ISCR during FY 2004. These activities, which are further detailed in the accompanying CD-ROM, fall

Table 3. ISCR's 38 participating guests for FY 2004.

Guest	Affiliation	LLNL Contact	Dates
Marian Brezina	University of Colorado	Rob Falgout	Dec. 1, 2000 – Nov. 30, 2004
Zhiqiang Cai	Purdue University	Charles Tong	Sept. 8, 2003 – Sept. 7, 2005
Praveen Chandramohan	Oak Ridge National Laboratory	Terence Critchlow	Nov. 7, 2003 – May. 6, 2004
Alok Choudhary	Northwestern University	Terence Critchlow	Sept. 15, 2003 – Sept. 13, 2005
Jennifer Dacles-Mariani	University of California, Davis	Garry Rodrigue	Sept. 1, 2003 – Aug. 30, 2005
Hans de Sterck	University of Colorado, Boulder	Rob Falgout	Oct. 1, 2003 – Sept. 30, 2004
Branden E. Fitelson	University of California, Berkeley	Terence Critchlow	Aug. 9, 2004 – Aug. 8, 2005
Franz Franchetti	Carnegie Mellon University	Kim Yates	June 4, 2004 – June 3, 2005
Alejandro Garcia	San Jose State University	Richard Hornung	Oct. 12, 2001 – Oct. 10, 2004
William Charles Gear	Princeton University	Steve Lee	Sept. 21, 2004 – Sept. 20, 2005
Matthew R. Gibbons	U.S. Airforce Academy	Bill Bateson	Aug. 9, 2004 – Aug. 8, 2005
Bernd Hamann	University of California, Davis	Mark Duchaineau	Aug. 9, 2004 – Aug. 8, 2005
Alan Hindmarsh	LLNL (retired)	Carol Woodward	Oct. 1, 2002 – Sept. 29, 2004
Martin Isenburg	University of North Carolina	Terence Critchlow	Nov. 4, 2003 – Nov. 3, 2004
Ken Joy	University of California, Davis	Mark Duchaineau	June 4, 2003 – June 3, 2005
Ramya Krishnamurthy	Oak Ridge National Laboratory	Terence Critchlow	Nov. 7, 2003 – Nov. 6, 2004
Johannes K. Kraus	University of Loeben	Van Henson	Nov. 30, 2002 – Sept. 28, 2004
Raytcho Lazarov	Texas A&M University	Panayot Vassilevski	Aug. 31, 2002 – Aug. 28, 2005
Oren Livne	Stanford University	Van Henson	Sept. 15, 2003 – Sept. 14, 2004
Bertram Ludaescher	San Diego Supercomputer Center	Terence Critchlow	Sept. 1, 2001 – Oct. 26, 2004
Jeannee Martin	LLNL (retired)	Bronis de Supinski	March 20, 2004 – March 19, 2005
Sally McKee	Cornell University	Bronis de Supinski	Oct. 14, 2001 – Oct. 13, 2004
Esmond Ng	Lawrence Berkeley National Laboratory	Edmond Chow	May 12, 2003 – May 11, 2004
Beth Ong	LLNL	Van Henson	Aug. 1, 2001 – Dec. 1, 2004
Peter Pacheco	University of San Francisco	Pat Miller	Aug. 9, 2004 – Aug. 8, 2005
Joseph E. Pasciak	Texas A&M University	Panayot Vassilevski	July 1, 2002 – June 29, 2005
Christoph Pflaum	Universität Erlangen-Nürnberg	Rob Falgout	Aug. 25, 2003 – Aug. 24, 2004
Elbridge Gerry Puckett	University of California, Davis	Louis Howell	June 30, 2003 – June 29, 2004
Markus Pueschel	Carnegie Mellon University	Kim Yates	April 14, 2003 – May 18, 2005
Ulrich Ruede	Universität Erlangen-Nürnberg	Rob Falgout	Aug. 15, 2000 – July 29, 2004
Paul E. Saylor	University of Illinois	Steve Lee	Aug. 22, 2004 – Aug. 21, 2005
Martin Schultz	Cornell University	Bronis de Supinski	March 1, 2003 – Dec. 30, 2004
M. Alex Schweitzer	Universität Bonn	Rob Falgout	Nov. 17, 2003 – July 22, 2005
Claudio Silva	University of Utah	Randy Frank	May 19, 2003 – June 3, 2004
Lansing Sloan	LLNL (retired)	Pete Eltgroth	June 16, 2004 – June 15, 2005
Christoph W. Ueberhuber	Techische Universität Wien	Kim Yates	March 1, 2003 – Feb. 28, 2004
Beata Winnicka	Argonne National Laboratory	Dan Quinlan	Aug. 15, 2003 – Aug. 14, 2004
Gabriel Wittum	Universität Heidelberg	Rob Falgout	Nov. 17, 2003 – Nov. 16, 2004

under two main themes: sponsored-research activities that stimulate interactions between academia and LLNL staff, and a diverse visitor program that enables both short- and long-term residential stays at LLNL.

ISCR oversees three different types of sponsored-research activities. The University Collaborative Research Program (UCRP), through the ISCR, funded seven research projects during FY 2004 at University of California campuses. These projects supported graduate students working on doctoral thesis research. The faculty principal investigators and students worked closely with an LLNL collaborator. The ISCR also coordinated the funding of 19 research subcontracts to various academic institutions throughout the United States. These contracts are normally funded by programs at LLNL to help address long-term Laboratory requirements. This type of vehicle is also used to fund sabbatical visits to LLNL for three to six months. Eleven faculty members spent at least a portion of their sabbatical leave here during FY 2004. With Laboratory Directed Research and Development (LDRD) funds, the ISCR also funds Exploratory Research in the Institutes (ERI). These research grants are awarded to LLNL staff with the goal of developing ties to academia through co-funded research projects. The ISCR oversaw three such projects in FY 2004. Annual summaries for LDRD projects, UCRP projects, and subcontracts can be found in the next three sections of this document.

Figure 1. ISCR Visitor Program FY1997–FY2004

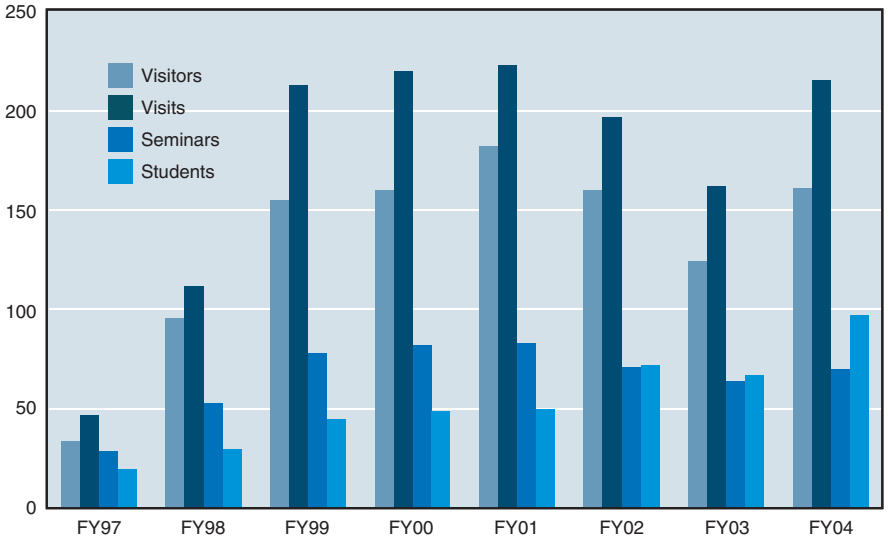

In FY 2004, the ISCR continued its tradition of an extensive and diverse Visitor Program. This program includes sabbatical visitors, sponsored workshops, summer students, and various seminars featuring external speakers. Altogether, the ISCR hosted 215 visits from 161 different visitors, an average of more than 4 visits per week. The vast majority (67%) of the visitors were from academia, with 9% from industry, 20% from other federal laboratories, and 4% from non-laboratory-based government institutions. Visitors from outside the United States made up 25% of the total.

The ASC Institute for Terascale Simulation Lecture Series was established in 2000 to enrich the intellectual atmosphere of LLNL's large simulation community through the visits of leaders representing the diverse areas of computation. In FY 2004, we hosted five speakers in this series. The general ISCR seminar series included an additional 65 talks covering a wide spectrum of research areas. Titles of all of these talks can be found in the Seminar Series section of this report, and associated abstracts can be found on the accompanying CD-ROM.

During the summer, ISCR hosted 94 visiting students for a total summer student population of 104, including the SEGRF students in residence year-round. The summer program exposes students to the stimulating and challenging work environment of a national laboratory. Successful candidates are hired as summer employees, assigned individual LLNL mentors, and given specific projects to which they will contribute. The project is appropriate to the student's background and skills, and ranges from programming tasks to original research. The topical coverage of the summer research program broadens each year as computation expands into new scientific areas and as computational tools become more powerful and diverse. Scalable algorithms, radiation transport, genomics, terascale visualization, and computer security are just a handful of topics from last summer's lively hallway conversations at the ISCR. The summer program runs from May to September, with most participants spending 10–12 weeks on site. Project

reports for most of the students can be found on the accompanying CD-ROM.

In June, with the advent of our large student summer program and sponsorship from the Defense Programs office of DOE Headquarters, we ramped up our fifth annual Summer Student Lecture Series—three different series on Computational Modeling at the Terascale, Computer Science at the Terascale, and Computer Security. Though the lectures were intended for students, permanent CASC researchers also attended. The 35 lecturers are listed in the Seminar section of this annual report.

Poster presentations were made by 41 ISCR summer students at the LLNL Student Research Symposium in August 2004. The event, held at LLNL's Central Café, attracted local media and scientific staff from across the Laboratory, as well as other students and summer research mentors. Overall, 139 students presented posters, making ISCR's share about 30% of LLNL's total. Topics of the posters ranged from Internet routing patterns to advanced gridding techniques for estuarine flow modeling; from performance modeling tools to wavelet-based compression of radiation opacities; from language interoperability to management of data in petabyte-scale file systems. Students ranged in seniority from community-college first-years to graduate students about to complete their doctorates.

Figure 1 charts the numbers of visitors and seminars over the past eight years. The number of students in residence in FY 2004 increased substantially due to the expansion of the ISCR's responsibility in LLNL's summer programs. CASC scientists mentored 55% of these students. Other LLNL organizations mentoring ISCR summer students were: CADSE, DCOM, EEBI, ICCD, NAIC, NIFE, and PAT. Some of these students elected to spend internships prescribed by their national fellowships at the ISCR, at no direct cost to the Laboratory, including DOE Computational Science Graduate Fellowship (CSGF) holders and Department of Homeland Security (DHS) fellows.

Finally, the ISCR sponsored or co-sponsored 10 scientific workshops in FY 2004. Two of these were hosted locally and exclusively by the ISCR; the rest were in cooperation with other organizations and held off-site. In each case, there is a vital LLNL interest and typically, several LLNL researchers participate. Reports on these workshops appear in a later section of this report.

Most of the raw material of this document comes directly from the visitors and principal investigators of their respective projects. We thank the Technical Information Department, including Alane Alchorn, Arnold Gatilao, Deanna Midtaune, and Al Miguel for their editorial work, Dan Moore for his graphic artistry in producing an easily navigable and visually pleasing document, and Maria Fogle and John Danielson for helping us finalize this document and getting it ready for print.

We hope that you enjoy examining this report on the ISCR's diverse activities in FY 2004. For further information about the Institute, please contact us at the address below. Inquiries about how you might enhance the ISCR program for FY 2005 and beyond are always welcome.

David E. Keyes

Institute for Scientific Computing Research
Lawrence Livermore National Laboratory
P. O. Box 808, L-561
Livermore, CA 94551
http://www.llnl.gov/iscr

David E. Keyes, Acting Director
http://www.llnl.gov/iscr/people/keyes
dekeyes@llnl.gov
925-422-1325

James R. McGraw, Deputy Director
http://www.llnl.gov/iscr/people/mcgraw
jmcgraw@llnl.gov
925-422-0541

Linda Becker, Administrator
http://www.llnl.gov/iscr/people/becker
becker22@llnl.gov
925-423-0421

Laboratory Directed Research and Development (LDRD)

The ISCR supported three Exploratory Research in the Institutes (ERI) projects with Laboratory Directed Research and Development (LDRD) funds during FY 2004. These research grants are awarded to LLNL staff with the goal of developing ties to academia through co-funded research projects. Anticipating the emergence of data science as a cross-disciplinary theme, the ISCR has concentrated its efforts in developing technologies for large-scale and distributed data sets for the past several years. The current portfolio contains a project on providing a single interface to diverse data sources on the Web, a project on tracking objects in a succession of images, and a project in multiresolution representation of scientific datasets for scalability across a range of computer architectures. This portfolio, originally motivated by purely scientific applications, has already paid dividends in some of the Laboratory's new homeland security applications.

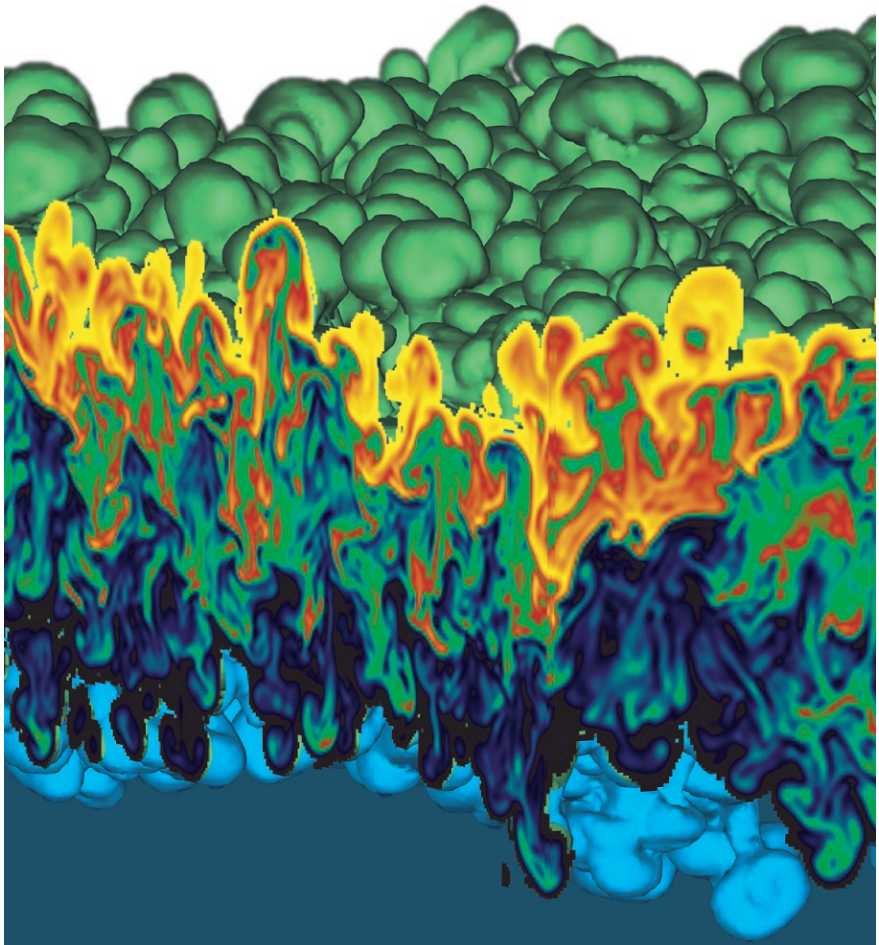

TOC

Enabling Large-Scale Data Access

Terence Critchlow and David Buttler, CASC

This project's goal was to develop an infrastructure capable of providing scientists with access to large numbers of data sources. To that end, we have developed an infrastructure that employs a user-specified description of a service of interest to crawl the Web. When an interface in the service class is identified, a wrapper that supports an XML-based query interface is automatically created. This year, we successfully demonstrated the application of our infrastructure on service classes from two vastly different domains — Basic Local Alignment Search Tool (BLAST) sequence similarity interfaces and interfaces to publication sites.

We have demonstrated the ability to automatically identify new data sources of interest while crawling the Web. We performed extensive testing of the infrastructure on hundreds of BLAST data sources and half a dozen publication sources. Next, we performed a minimal Web crawl, during which we identified and wrapped nine previously unknown BLAST sites and one new publication site. In addition, we identified approximately 100 Web pages that contained citations but were not queryable.

This work supports national security and other LLNL missions by benefiting ongoing programs at LLNL, such as Department of Homeland Security nonproliferation and detection efforts, that must utilize information from a wide variety of sources, including some that cannot be easily integrated using traditional techniques. Our infrastructure simplifies the process of creating an interface that combines local data with related information publicly available over the Internet, such as scientific publications.

We met and exceeded our FY 2004 milestones. In addition to demonstrating end-to-end automatic wrapper generation for BLAST interfaces, we extended the service class description to handle more complex data types, such as citations that cannot be easily represented by regular expressions. We performed several short Web crawls that successfully demonstrated the application of our infrastructure on both well-structured BLAST interfaces and poorly structured publication interfaces by identifying previously unknown sites in both domains. The expertise gained in this project will be applied to the Biodefense Knowledge Center Bio-Encyclopedia effort.

Publications

(1) D. Rocco, T. Critchlow. "Automatic Discovery and Classification of Bioinformatics Web Sources." *In Proceedings of the Georgia Tech Conference on Bioinformatics*. UCRL-JC-152980. 2003

(2) J. Caverlee, L. Liu, D. Buttler. "Probe, Cluster, and Discover:focused extraction of QA-Pagelets from the deep web." *In proceedings of the IEEE conference on Data Engineering*. 2004

(3) Wei Han. Ph.D. dissertation Georgia Institute of Technology 2003

(4) Dan Rocco. Ph.D. Dissertation, Georgia Institute of Technology 2004

Detection and Tracking in Video

Principal Investigator

Chandrika Kamath, CASC

Video cameras are used for monitoring and surveillance in several applications. We are developing robust, accurate, and near-real-time techniques for detecting and tracking moving objects in video from a stationary camera. This allows us to model the interactions among the objects, thereby enabling us to identify normal patterns and detect unusual events. Our algorithms and software include techniques to separate the moving foreground from the background, extract features representing the foreground objects, track these objects from frame to frame, and post-process the tracks for display. We focus on video taken under less-than-ideal conditions, with objects of different sizes moving at different speeds, occlusions, changing illumination, low resolution, and low frame rates.

The capability to detect and track in video supports the national security mission of LLNL by enabling new monitoring and surveillance applications for counterterrorism and counter-proliferation. This project will produce robust and accurate technology for video detection and tracking under less-than-ideal conditions with occlusions, fog or changing illumination, or at a low resolution or frame rate. This project will enhance existing algorithms to address these situations, allowing us to better understand their limitations, which in turn, will determine the conditions under which successful surveillance is possible. The algorithms and software are being applied to surveillance video, as well as spatiotemporal data from computer simulations.

During FY 2004, we

(1) Created a software infrastructure to handle streaming video data.

(2) Implemented several background subtraction algorithms and evaluated them on videos taken under different conditions.

(3) Proposed a new background subtraction method that outperforms other methods, especially on low-resolution, low-frame-rate video.

(4) Extracted features and used them in simple tracking algorithms.

We also filed a provisional patent on the new method and summarized our work in two papers. We are currently adapting the tracking algorithms to work under adverse conditions. We collaborated with the University of Colorado, Boulder, on tracking people, the University of California, San Diego, on tracking under occlusions, and a summer student on object representations for tracking.

Publications

Cheung, S.C., C. Kamath, "Robust background subtraction with foreground validation for urban traffic video," *Eurasip Journal on Applied Signal Processing*, UCRL-JRNL-201916.

Moelich, M., "Autonomous motion segmentation of multiple objects in low resolution video using variational level sets," UCRL-TR-201054.

Gyaourova, A., C. Kamath, S.C. Cheung, "Block matching for object tracking," UCRL-TR-200271.

Cheung, S.C., C. Kamath, "Robust techniques for background subtraction in urban traffic video," *Video Communications and Image Processing Conference*, **Vol. 5308**, pp. 881-892, UCRL-CONF-200706.

Cheung, S.C., "Robust techniques for background subtraction," UCRL-ABS-200371.

Our new technique for background subtraction (bottom) is less sensitive to changes in illumination in comparison with current techniques (top). The pixels highlighted in purple indicate the moving objects in the frame.

ViSUS: Visualization Streams for Ultimate Scalability

Principal Investigator

Valerio Pascucci, CASC

We are developing a suite of progressive visualization algorithms and a streaming infrastructure to enable interactive exploration of large scientific data sets. The methodology optimizes the data flow in a pipeline of processing modules. Each module reads and writes a multi-resolution representation of a geometric model, providing the flexibility to trade speed for accuracy, as needed. The data flow is streamlined with progressive algorithms that map local geometric updates of the input into immediate updates of the output. A prototype streaming infrastructure will demonstrate the flexibility and scalability of this approach for visualizing large data sets on a single desktop computer, a cluster of personal computers, and heterogeneous computing resources.

In FY 2004, we brought ViSUS to a level of maturity and robustness allowing direct deployment for a number of targeted users. The main milestones achieved include developing new techniques that accelerate isosurface extraction with occlusion culling, graphics hardware, and view-dependent refinements. A test viewer has been developed for datasets from the HYDRA simulation. We provided a stable library with full implementation of our streaming technology that can be used by simulation codes for saving rectilinear grids in ViSUS IDX format. We released a new version of the Progressive Viewer with full slicing, isocontouring, and volume-rendering capabilities. We are working in collaboration with the MIRANDA team to start using the IDX format as output of choice for the Blue Gene/L runs.

The ViSUS project benefits the Laboratory at least at two levels. At the deployment level, the improved efficiency in the use of hardware resources reduces the cost of visualization-hardware infrastructures. At the scientific level, the developed technology reduces the overall time required for the design, simulation, and visualization cycle. The ability to remotely monitor large and expensive simulations saves computing resources through early termination and restart of erroneous test simulations. Runtime steering will be possible for simulation codes with mechanisms for dynamic modification of running conditions.

Use of our innovative, high-performance visualization techniques allows interactive display of very large data sets on simple desktop workstations and the monitoring (or steering) of large parallel simulations. This will have valuable applications to several LLNL missions, including stockpile stewardship, nonproliferation, energy security, and environmental management, that use large-scale modeling and simulations.

Publications

Laney, D. and V. Pascucci, "Progressive Compression of Volumetric Subdivision Meshes," *Proceedings of the International Symposium on 3D Data Processing, Visualization, and Transmission* 2004, pp. 293–300, UCRL-CONF-203679.

Pascucci, V. "Topology Diagrams in Scientific Visualization" Chapter in: *Surface Topological Data Structures: An Introduction for Geographical Information Science*, pp. 121–130, UCRL-200013-BOOK.

Pascucci, V., "Isosurface computation made simple: Hardware acceleration, adaptive refinement and tetrahedral stripping," *Proceedings of the Joint Eurographics - IEEE TVCG Symposium on Visualization*, pp. 293–300, UCRL-CONF-202459.

Van Kreveld, M.J., R.van Oostrum, C.L. Bajaj, V. Pascucci, and D.R. Schikore, "Efficient contour tree and minimum seed set construction" Chapter in: *Surface Topological Data Structures: An Introduction for Geographical Information Science*, pp. 71–86, UCRL-200018-BOOK.

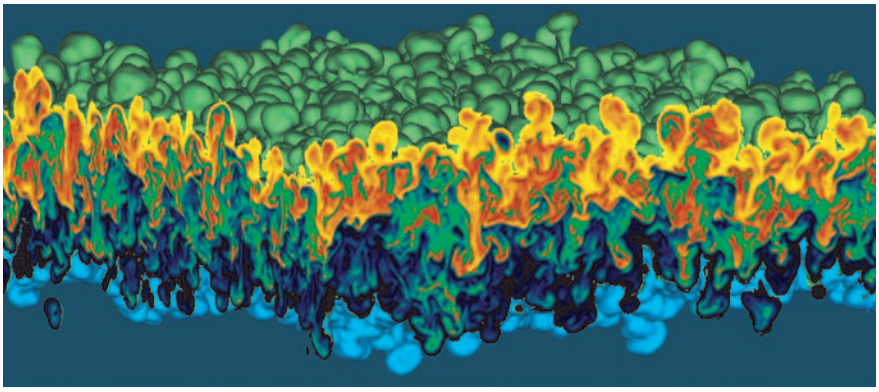

The ViSUS Progressive Viewer generates a snapshot of the MIRANDA hydrodynamics code computing a Rayleigh–Taylor turbulent mixing of fluids.

TOC

University Collaborative Research Program

The ISCR supported seven University Collaborative Research Program (UCRP) projects during FY 2004 at the University of California campuses. These projects support graduate students working on doctoral thesis research, and the faculty principal investigators and students work closely with an LLNL collaborator. The projects in FY 2004 spanned four different UC campuses, as well as the departments of civil engineering, computer science, mechanical engineering, and physics. They ranged from understanding and optimizing the performance of advanced architecture computers to improving engineering simulation models; from the development of tracking and recognizing objects to basic computational science at the frontier.

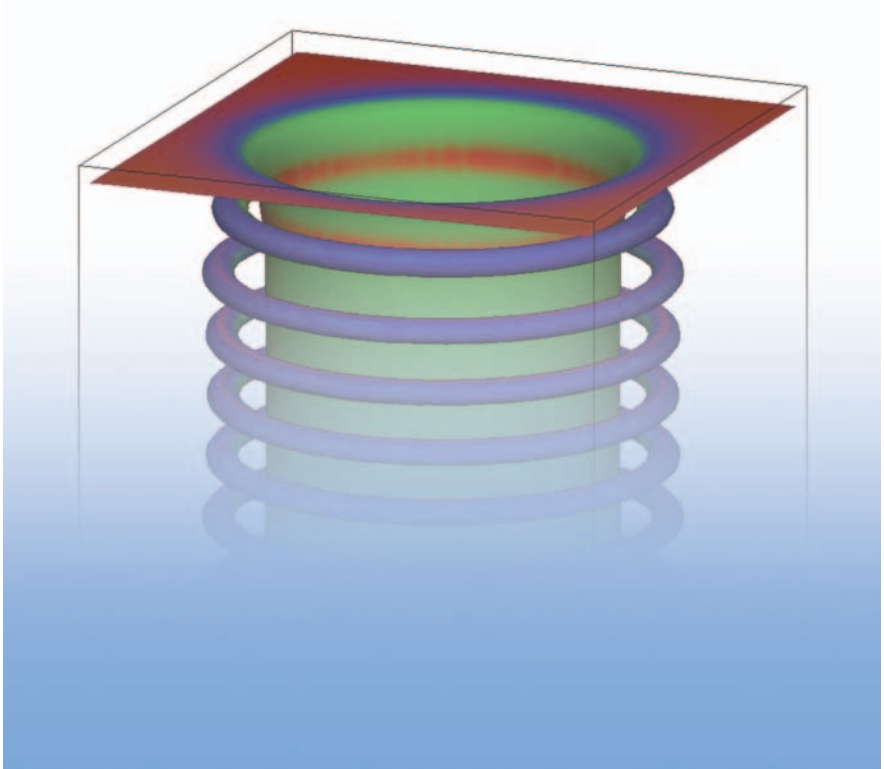

TOC

Data-Driven Execution of Communication-Tolerant Algorithms

Principal Investigator **Scott B. Baden,** University of California, San Diego

Collaborator **Daniel Quinlan,** CASC

Communication-tolerant algorithms are expected to play an important role in achieving scalability on large-scale platforms with many thousands of processors. However, communication-tolerant application development is time-consuming and prone to error, even for the expert programmer. A run-time substrate that supports non-bulk, synchronous, data-driven execution is under investigation. Because it treats communication and computation as coupled activities, rather than as distinct phases of execution, the data-driven execution model naturally supports irregular or unpredictable communication delays that are expected on large-scale platforms, such as BlueGene/L.

The data-driven substrate programmer expresses an application in terms of partially ordered communication and computation operations, as constrained by the data dependences inherent to the application. The substrate invokes an externally specified scheduler to determine the precise task-execution ordering. The user may further constrain the scheduling by means of *performance metadata* decorating the graph. The user expresses partially ordered operations in terms of a directed graph called a *task graph*. Vertices in the task graph correspond to tasks and edges to data dependences between the tasks.

The substrate is implemented as a C++ library and runs as a background thread (i.e., a proxy) to manage dependence information and task scheduling. This background thread, called the Mover–Dispatcher, routes data from completed tasks to dependent tasks and determines when such tasks are enabled for execution.

When several possible tasks are ready for execution, their exact order is determined by the scheduler in conjunction with performance metadata. We built the communication proxy part of the Mover–Dispatcher, which supports non-blocking, asynchronous communication by means of multithreading. We are currently testing this capability on a latency-tolerant variant of a Gauss Seidel red-black Poisson solver in three dimensions using a Beowulf cluster in our research lab.

The interpretation of performance metadata is up to the application. For example, it might express priorities or affinities. We have tested out this capability on blocked LU running on a single processor. Our goal was to reorder computations in order to improve memory locality. The blocked LU application demonstrated the importance of an economical run-time representation for the task graph structure. The overhead of dynamically reordering task execution is 25% of the overall running time and an improved run-time representation for tasks graph representation is under investigation. We are exploring two optimizations. The first takes advantage of graph sparsity—most nodes depend on only a few other nodes. The second takes advantage of temporal locality. Rather than instantiate an entire task graph at initialization time, another approach is to create task graph nodes on demand and to recycle disused nodes. This strategy effectively throttles task generation and avoids tying up memory with task information that won't be needed for a long time.

Publications

S. B. Baden, "Moving forward in large scale computation" in J Dongarra, K. Madsen, J. Wasniewski (Eds.), 7th International Conference on Applied Parallel Computing (PARA '04), Lyngby Denmark, June 2004, *Lecture Notes in Computer Science,* Springer, 2004.

S. B. Baden, "Masking Latency with Data Driven Program Variants" in J Dongarra, K. Madsen, J. Wasniewski (Eds.), 7th International Conference on Applied Parallel Computing (PARA '04), Lyngby Denmark, June 2004, *Lecture Notes in Computer Science,* Springer, 2004.

Feature-Based Approaches for Long-Range Motion Segmentation and Object Tracking

Principal Investigator **Serge Belongie,** University of California, Davis

Collaborator **Chandrika Kamath,** CASC

Image segmentation is the problem of partitioning the pixels in an image into a relatively small number of regions that correspond to objects or parts of objects. It is one of the hardest (and oldest) open problems in computer vision, and it plays an important role in the process of object detection and recognition. As challenging and computationally intensive as image segmentation is, it also happens to be a problem that the human visual system solves effortlessly. The goal of this project is to develop methods for image and video segmentation with an emphasis on motion-based processing.

Our proposed work divides into two main areas: "what went where" and "who went where." The first area, on which we have already begun working, addresses the problem of motion segmentation for image sequences with large inter-frame displacement, e.g., more than 10% of the image width. The second area, for which the principal application area will be people tracking, deals with the problem of detecting and tracking moving objects that exhibit temporally periodic variation in appearance.

We first describe our solution to the "what went where" problem. The algorithm operates in two stages, starting with robust estimation of the underlying motion fields and concluding with dense assignment of pixels to motion fields. The first stage of this process is the first dense-motion segmentation method to operationalize the layer-based formulation for multiple discrete motions. We detect the motion layers using a variant of RANSAC (Random Sample Consensus) on detected interest points using a planar projective (2D homography) motion model. We perform the layer assignment using a fast, graph-cut-based MRF (Markov Random Field) formulation, which enforces spatially piecewise smooth-pixel assignments.

The case of moving objects that do not fit the planar projective model is addressed in the second stage. A poor fit to the planar projective model can arise from two main causes: non-planarity and non-rigidity. For this purpose, we refine the fit via iterative regression using a regularized thin-plate spline (TPS).

Our approach to the "who went where" problem builds on our approach to the first problem by interfacing it in a novel way with the classical tools of multiview geometry. In particular, we apply the knowledge that sets of frames depicting a periodically moving object in a given shape but with varying poses will (approximately) satisfy the multiview constraints for a rigid object. After automatically detecting the period of a given object in the field of view, we apply the "what went where" framework to segment the object from the background. The segmented foreground object across the period-separated frames is then treated as a conventional stereo- or multiview-reconstruction problem. Using this framework, we demonstrated a 3D reconstruction of sparse interest points on a video sequence of a pedestrian and showed preliminary results computing dense disparity using a graph-cut-based stereo correspondence engine.

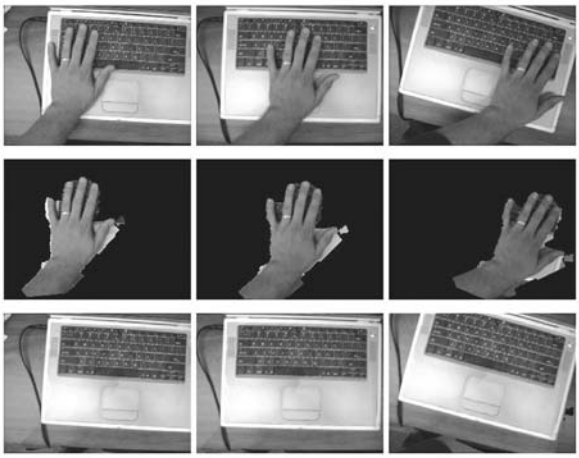

Illustration of video object deletion. Original frames (top). Segmented layer corresponding to hand motion (middle). Reconstruction without the hand layer using the recovered motion of the keyboard (bottom).
Note that no additional frames beyond the three shown were used as input.

Publications

Sameer Agarwal, Satya Mallick, David Kriegman and Serge Belongie, "On Refractive Optical Flow," *European Conference on Computer Vision*, Prague, Czech Republic, pp. 483–494, **vol. 2**, 2004.

Josh Wills and Serge Belongie, "A feature based method for determining dense long range correspondences," *European Conference on Computer Vision*, Prague, Czech Republic, pp. 170–182, **vol. 3**, 2004.

Serge Belongie and Josh Wills, "Structure from Periodic Motion," *International Workshop on Spatial Coherence for Visual Motion Analysis*, Prague, Czech Republic, 2004.

DNS and Modeling of Dispersion of Solid or Liquid Particles in Turbulent Flows

Principal Investigator **Said Elghobashi,** University of California, Irvine

Collaborator **Robert Lee,** Atmospheric Science Division

The release and subsequent transport of toxic materials into the atmosphere or inside and around buildings are of major concern to national security. These toxic materials can be in the form of gas, liquid droplets or solid particles and often undergo chemical reactions during the transport process. An important distinguishing feature of the transport of toxic materials into the atmosphere or inside and around buildings is the wide spectra of length scales and time scales involved. Length scales range from microns (particle size) to kilometers (turbulent eddy size) and the corresponding time scales range from microseconds to days. The approach of the proposed research is to employ direct numerical simulation (DNS) for prototypical particle-laden turbulent flows (e.g., flow over a backward-facing step) to evaluate the turbulent *correlations directly* and provide accurate models for them that can be introduced in LES (Large Eddy Simulation) or RANS (Reynolds-Averaged Navier–Stokes equations) codes.

The proposed work encompasses the numerical solution of the three-dimensional, time-dependent Navier–Stokes and continuity equations of the turbulent flow in addition to solving the equations of motion of the dispersed particles. The time-averaged correlations needed for RANS equations will be evaluated from their instantaneous values in DNS via ensemble and time averaging, and then models will be developed to relate these correlations to the dependent variables used in RANS equations.

Our work on the project during the past six months concentrated on the mathematical development of the necessary boundary conditions for the DNS code for the backward-facing step and performing short runs on a small parallel computer. We also spent considerable effort in converting our DNS parallel code that was running on a Cray-T3E so that it can run on a MCR computer (LLNL). A new machine-independent Fast Fourier Transform (FFT) was written and incorporated into the code. We are starting the DNS of the flow over the backward-facing step.

Numerical Study of Coexisting Superconductivity and Ferromagnetism: Applications to Real Materials

Principal Investigator

Warren E. Pickett, University of California, Davis

Collaborator

Francois Gygi, CASC

The discovery of ferromagnetic metals that become superconducting at lower temperatures (UGe_2, $ZrZn_2$, URhGe) early in this young century provided one of the big surprises in materials physics in recent years. For more than four decades, such coexistence was believed to be all but impossible. This coexistence of two competing phases, each a macroscopic manifestation of the quantum behavior of electrons, fully qualifies this as a new state of matter. The goal of this project is to include material-specific characteristics of the metallic states in the underlying formalism and to perform numerical studies to illuminate the microscopic driving mechanisms.

Our approach consists of three prongs:

(1) Reformulating the theory at its more basic levels to gain insight

(2) Adapting the formalism, developing algorithms and making the resulting codes applicable to real materials, such as those mentioned above, as opposed to the earlier study of model systems

(3) Computing solutions to the superconducting gap equations to map out phase diagrams.

Our work has moved to simulating the novel type of new coexisting ferromagnetic-superconducting state called FFLO, after its discoverers Fulde, Farrell, Larkin and Ovchinnikov. This state, which can occur in very weak ferromagnets or for paramagnets near the critical value of magnetic field, consists of the coalescence of superconducting pairs of electrons with a non-vanishing pair momentum. The resulting state and its order parameter are inhomogeneous in space. Our primary results are the identification of phase boundaries that specify the phase diagram.

The results of this year's research are:

(1) Identification of a characteristic velocity spectral function that is a fundamental part of the integral of the integral equation

(2) Development and implementation of an algorithm to evaluate this function

(3) Showing how its spectrum affects the phase diagram and obtain predictions for a specific candidate, $ZrZn_2$.

During the past year, two new materials have been found to display some of the characteristics of the FFLO phase. These two systems are the peculiar rare-earth-based superconductor $CeCoIn_5$ and one of the first "heavy fermion" superconductors, UBe_{13}. If confirmed to be FFLO, these will add to the single system (the two-dimensional organic metal k-$(BEDT\text{-}TTF)_2Cu(NCS)_2$) that is relatively well established to enter an FFLO phase when a magnetic field is applied. One strong focus of the coming year will be to apply our new methods to these novel superconductors and help assess whether an FFLO phase really does arise. The computational facilities of ISCR are anticipated to be important for the success of the latter stages of this project.

Publications

A. B. Kyker, W. E. Pickett, and F. Gygi, "Fermiology and Fulde-Farrell-Larkin-Ovchinnikov Phase Formation," *Physical Review B* (submitted), 2004.

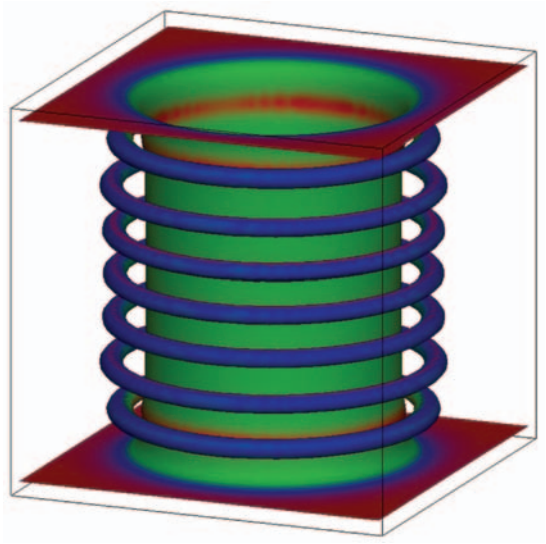

A current-carrying wire penetrates a slab of superconductor where it is electrically insulated, producing a radial magnetic field. Such a field, penetrating a type-II superconductor, produces regions (flux lines) of depressed superconducting strength, each containing one quantum of magnetic flux. For this geometry, the flux lines align regularly in rings, a generalization of the Abrikosov vortex lattice state in superconducting films. The color here gives the rate of change of the superconducting order parameter on the various isosurfaces (red is slow, blue is rapid).

Lagrangian Simulation of Penetration and Other Extreme-Deformation Events: Moving Beyond Meshless Methods

Principal Investigator **Mark Rashid,** University of California, Davis

Collaborator **Mike Puso,** Engineering

Solid mechanics problems involving extreme-deformation events, such as projectile/target interaction and fragmentation of cased explosives, present a considerable challenge to the conventional Lagrangian finite element method (FEM). Accordingly, beginning in the mid-1990s, various so-called meshless approximation methods have been proposed as possible alternative approaches. This research project involves the development of a variational approximation method that attempts to combine the best features of both the conventional FEM and meshless methods.

The technical approach involves a synergy of two distinct and innovative elements. In the first of these, the solid-mechanics problem is discretized using a flexible geometric subdivision scheme in which "elements" take the form not of hexahedra or tetrahedra but of arbitrary polyhedra. This flexibility vastly simplifies the task of mesh generation and thereby facilitates periodic remeshing, as is often required by extreme mesh distortion. The second element of the project involves a novel method for transferring information relating to the current material state from the old, distorted mesh to the new mesh following a remeshing cycle. This *remapping* step is required before the solution can be continued with the new mesh and can be a significant source of errors.

VETFEM is a finite-element-like Galerkin approximation method in which the basis functions and numerical quadrature are facilitated by a polyhedral partition of the material domain. In contrast to the conventional FEM, wherein all elements must conform to strict geometric and topological requirements, VET elements can take arbitrary polyhedral forms. The VET element formulation has been successfully extended to 3D, coded, and integrated into an existing research code. This code offers flexible and expandable capabilities, including explicit dynamics, implicit dynamics, quasistatics, strongly-objective (Rashid, 1993) finite-deformation one.

As with the VET element formulation, the extension of the existing 2D state remapping methodology to 3D turned out to be challenging. The fundamental algorithmic problem associated with the variational remapping method is that of partitioning the volume of a "new" tributary region among a tiling of "old" tributary regions with which it overlaps. The key innovation involves a highly efficient and robust optimization procedure that closely approximates the required volume partition without resorting to any cumbersome geometric calculations. This has been accomplished and the result is a highly efficient and robust piece of code.

A preliminary result is illustrated in figure, which shows a twisted toroidal mesh consisting of eight-node hex elements. A fairly complex function was defined on this mesh, and then it was rotated about its axis through 90 degrees in a series of steps. After each increment of rotation, the integration-point values of the function were remapped from the "old" mesh to the new one. Because the function itself possessed a periodicity of 90 degrees, the final, succession can be directly compared to the original ones. The error for different numbers of rotation increments is shown in the table. The performance of the new remapping method is being studied with this and other types of analyses.

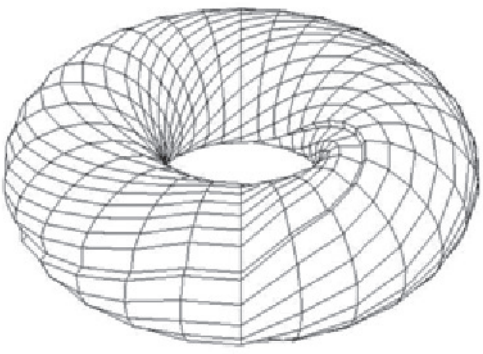

Twisted toroidal mesh is rotated 90 degrees in a series of steps, with remapping occurring after each increment of rotation. The error is shown for a range of the number of steps.

No. of rotation increments	RMS error	Max error
1	0.013404	0.027977
2	0.0132459	0.0441313
3	0.0148609	0.0547436
4	0.0174637	0.0616309
5	0.0193817	0.0616309
6	0.0212833	0.07386956
7	0.0228581	0.0830038
8	0.0244043	0.087037
9	0.02591708	0.1009022
10	0.027899	0.166627

Publications

M.M. Rashid and M. Selimotic "A Three-Dimensional Finite Element Method With Arbitrary Polyhedral Elements," submitted, 2004

Memory Access Pattern Signatures and Certificates of Relevance for Benchmarks

Principal Investigator

Allan Snavely, San Diego Supercomputer Center and University of California, San Diego

Collaborator

Bronis de Supinski, CASC

This award has supported Michael O. McCracken in his fundamental research investigations into performance modeling and prediction as part of a collaboration involving the University of California, San Diego's Computer Science Department (where McCracken is a student and Allan Snavely an adjunct assistant professor), the San Diego Supercomputer Center's (SDSC's) Performance Modeling and Characterization (PMaC) laboratory led by Snavely, and LLNL's Bronis de Supinski. Michael has thus been able to participate in research relevant to both organizations and also to PERC (the Performance Evaluation Research Center), a DOE Office of Science ISIC.

So far our research has investigated the performance implications of memory access patterns and useful definitions of "signature distance" between the memory access patterns of different basic blocks from the same or different programs. The goal is to improve the accuracy and speed of the Convolution method for performance prediction. We have enhanced the functionality of the MetaSim tool for gathering memory access pattern signatures and have made this tool platform independent. We have been investigating what kinds of memory access patterns exist "in nature" and exploring the performance implications of memory access patterns. We have developed a nomenclature and symbolic representation of memory access patterns leveraging previous work by Nick Mitchell.

We have been examining definitions of "signature distance" between basic blocks—our expectation is that basic blocks with similar memory access patterns will perform similarly on a given machine. In work defining a meaningful metric for "signature distance," he established orthogonal properties of loops including memory footprint, memory access pattern, type and intensity of floating-point, and ILP (instruction level parallelism) operations that could provide "certificates of relevance" for benchmarks. This work has demonstrated that reasonably accurate estimates of an application's performance can be derived from its MetaSim profile and mapping its basic blocks to a similar set of benchmark basic blocks with an established performance profile. McCracken has published a paper showing how these performance model predictions can guide dynamic algorithm selection.

Currently McCracken is developing a study of the HPC Challenge Benchmarks that will position these, and several strategic DOE applications, in dimensions of spatial and temporal locality. This work was presented at the HPCS meeting at SC04 in Pittsburgh. Further direct collaboration with LLNL is resulting in the means to acquire the application profiles via static analysis using the ROSE compiler infrastructure. This approach can result in significant time-savings and is leveraging de Supinski's expertise in compiler technology. The idea is that much of the information currently acquired via tracing, such as operation types and counts, communication and memory access patterns, can be determined or at least reasonably estimated more rapidly via static analysis. Tracing will then be used simply to confirm compile-time information via sampling and to fill in information (such as loop bounds) that may in some cases be unknown at compile time. The goal is to speed up code profiling by an order of magnitude.

Horizon Machine Profile with MetaSim data.

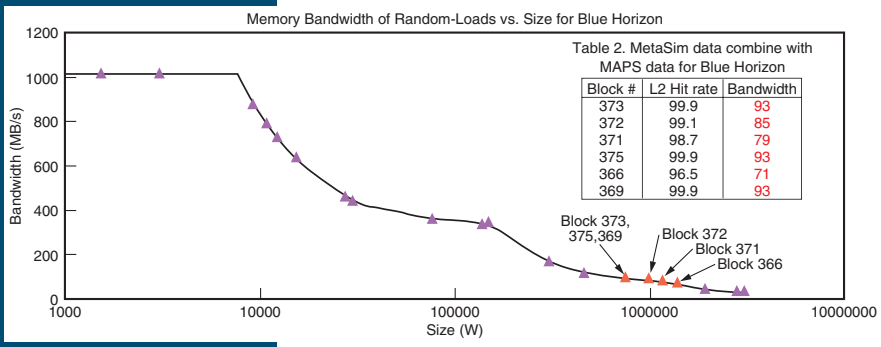

Table 2. MetaSim data combine with MAPS data for Blue Horizon

Block #	L2 Hit rate	Bandwidth
373	99.9	93
372	99.1	85
371	98.7	79
375	99.9	93
366	96.5	71
369	99.9	93

Publications

M.O. McCracken, A. Snavely, A. Malony, "Performance Modeling for Dynamic Algorithm Selection," *ICCS Workshop on Performance Modeling and Analysis (PMA03)*, June 2003, Melbourne, Australia.

TOC

Visual Tracking and Recognition for Biometrics and Interactive Visualization

Principal Investigator **Matthew Turk,** University of California, Santa Barbara

Collaborator **Lenny Tsap**, Electronics Engineering Technologies

At the Four Eyes Lab in the Computer Science Department at the University of California, Santa Barbara, we have pursued research on visual tracking and recognition for biometrics and interactive visualization with notable progress and success. Computer vision is a promising and powerful sensing modality that can be used to unobtrusively track, model, and classify human appearance and behavior.

Our primary goals for the year were to make progress in

1. Face tracking and facial expression analysis
2. Hand detection, tracking, and gesture recognition
3. Extracting and using reliable depth edges in images and video.

These areas contribute to the overall goals of the project and give us a solid start in addressing the technical needs of biometrics and visualization problems. In face-tracking and facial-expression analysis, we extended our wavelet-based method to enable real-time tracking, and we used embedding techniques to represent high-dimensional face information in low-dimensional manifolds to investigate the representation and recognition of dynamic facial expressions [1, 2]. We implemented and tested real-time hand tracking and recognition and successfully applied it to mobile computing applications [3,4], winning a workshop Best Paper award. We also developed a technique for reliably extracting depth edges using a digital camera with multiple flashes and developed methods to produce non-photorealistic renderings of scenes and recognize "fingerspelling" letters (used in sign language) from images, which resulted in additional publications [5,6].

In the past decade, there has been an increasing interest in moving beyond the traditional applications of computer vision, such as robot navigation, object recognition, and industrial inspection, to using vision technology as an effective input modality in human-computer interaction (HCI) often pursued in the context of multimodal or perceptual interfaces. The general focus of these efforts is to integrate multiple perceptual modalities, such as computer vision, speech and sound processing, and haptic I/O, into the user interface. Such video-based sensing is passive and non-intrusive as it does not require contact with the user or any special-purpose devices.

The primary tasks of computer vision in these scenarios are to detect, model, recognize, and interpret various visual aspects of human behavior. If delivered reliably and robustly, such vision technology can support a range of functionality in

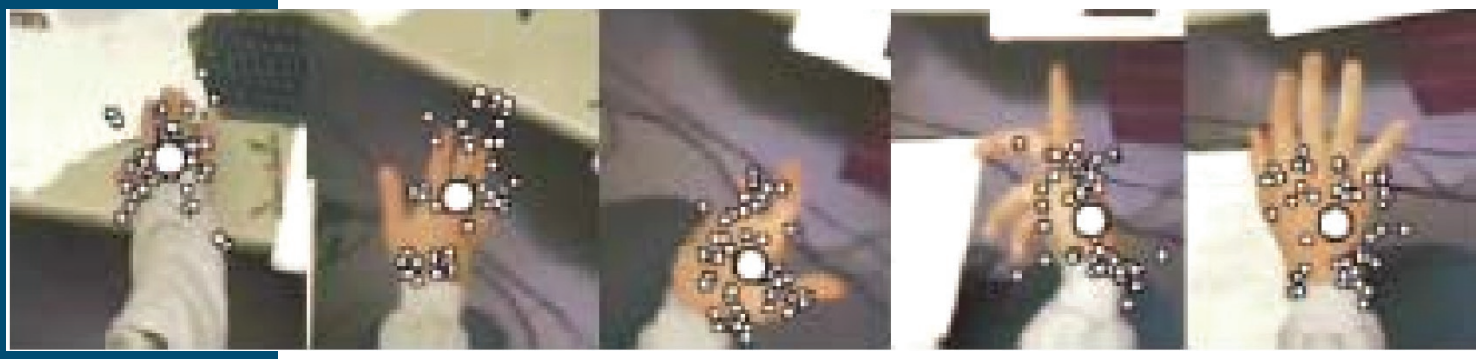

Hand tracking using the "Flocks of Features" method.

interactive systems by conveying relevant visual information about the user, such as identity, location, and movement, thus providing key contextual information. In order to fully support visual aspects of interaction, several tasks must be addressed, such as face detection and recognition, head and face tracking, facial expression analysis, eye-gaze tracking, body tracking, hand tracking, gait recognition, and recognition of postures, gestures, and overall activity.

Current biometric systems are based on processing images of a user's face, iris pattern, hand geometry, fingerprint, and other physical characteristics. As dynamic tracking and recognition technologies mature, behavior may well become an important aspect of biometrics – e.g., how people make facial expressions, how they move, their gaze and blinking patterns, and their postures. Leveraging these technologies will require a different approach to biometrics, with integration across modalities and across time (temporal integration) becoming central to the problem of verification.

Publications

1. C. Hu, Y. Chang, R. Feris, M. Turk, "Manifold based analysis of facial expression," *IEEE Workshop on Face Processing in Video*, Washington, D.C., June 2004.

2. Y. Chang, C. Hu, M. Turk, "Probabilistic expression analysis on manifolds," *International Conference on Computer Vision and Pattern Recognition*, Washington DC, June 2004.

3. M. Kölsch and M. Turk, "Fast 2D hand tracking with flocks of features and multi-cue integration," *IEEE Workshop on Real-Time Vision for Human-Computer Interaction*, Washington DC, USA, June 2004 (Best Paper Award).

4. M. Kölsch and M. Turk, "Analysis of rotational robustness of hand detection with rectangle features," *International Conference on Pattern Recognition*, Cambridge, U.K., August 2004.

5. R. Feris, M. Turk, R. Raskar, K. Tan and G. Ohashi, "Exploiting depth discontinuities for vision-based fingerspelling recognition," *IEEE Workshop on Real-Time Vision for Human-Computer Interaction*, Washington DC, USA, June 2004.

6. R. Raskar, K. Tan, R. Feris, J. Yu, M. Turk, "A non-photorealistic camera: depth edge detection and stylized rendering using multi-flash imaging," *ACM SIGGRAPH*, Los Angeles, August 2004.

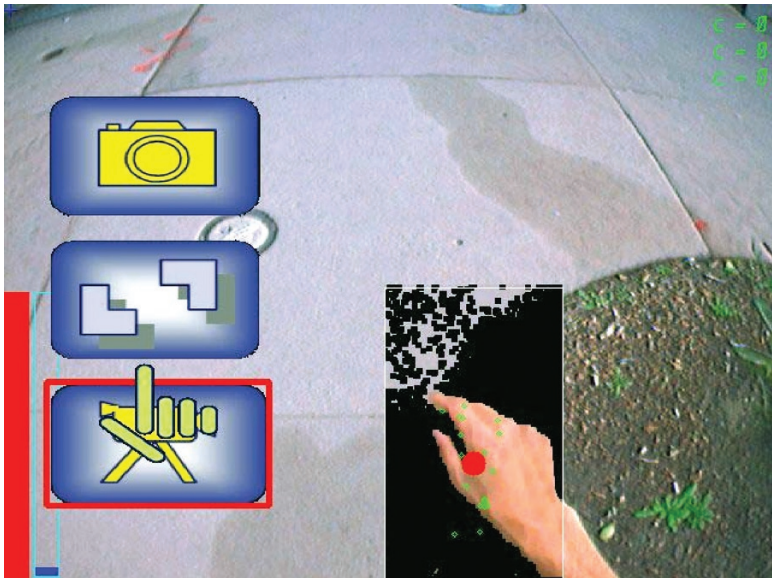

Figure 2. A user interface employing hand gestures.

Research Subcontracts

The ISCR supported 19 research subcontracts to various institutions throughout the United States. These contracts are normally funded by programs at LLNL to help address long-term Laboratory requirements. These subcontracts typically fund residential visits by university faculty for close collaborations with scientists in the Computation Directorate. Brief reports follow detailing ongoing work enabled by these subcontracts.

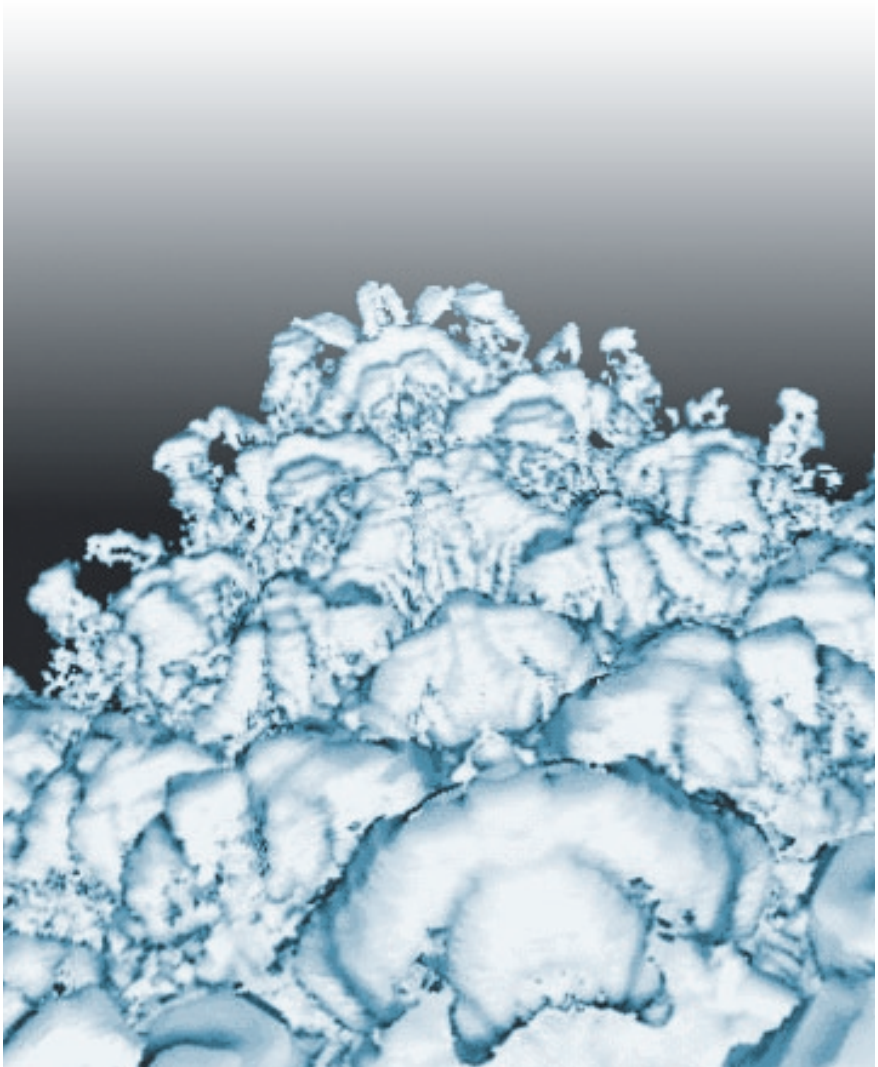

TOC

Nanohydrodynamic Simulation of the Rayleigh–Taylor Instability

Principal Investigator

Berni Alder, University of California, Davis

Collaborator

Garry Rodrigue, CASC

During the past year, we collaborated with a group at Los Alamos National Laboratory (LANL) and successfully carried out molecular dynamics simulations of the Rayleigh–Taylor instability that quantitatively predicts both the early-time behavior given by linear stability analyses and the long-time behavior given by experiment, as well as some of the numerical solutions of the continuum Navier–Stokes equation. This was published in *Proceedings of the National Academy of Sciences* and helped establish why different long-time behavior was observed in various continuum calculations.

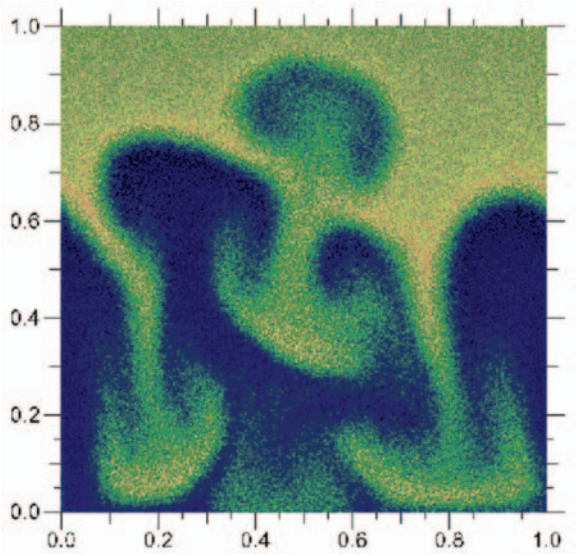

Demonstration of the Rayleigh–Taylor instability. The color variation goes from yellow at the top for the highest density to dark blue for the lowest density.

This is the first time that three-dimensional particle methods quantitatively simulated turbulent mixing. In the biggest and only molecular dynamics run, a qualitatively different asymptotic behavior was found, namely the heavy fluid formed drops instead of coalescing into a single slug. This could be real and appear only in particle simulations due to the presence of fluctuations, which could cut the thin thread by which the mushrooms are connected to the tip of the spike.

In order to firmly establish this phenomenon, which is of great importance to groups such as the NIF and because the mixing region advances more slowly (i.e., time to the first instead of the second power), we have developed an alternative particle algorithm that is 100 times faster and allows us to study even bigger systems for longer times. This was accepted for publication in *Molecular Physics*. To show the power of this scheme, we have attached a graph of the Rayleigh–Taylor instability at fairly long times with a few million particles that were produced on a single–processor computer in a few hours.

What these preliminary results indicate is the necessity to confirm them by carrying out the biggest Rayleigh–Taylor calculation possible on the IBM BlueGene/L machine for which the new particle algorithm is particularly well suited. We also need to carry out a comparison to the best continuum solution to see the quantitative effects of the nonlinear terms and fluctuations left out in the Navier–Stokes equations.

TOC

Multiple Animal Tracking in the Smart Vivarium

Principal Investigator **Serge Belongie,** University of California, San Diego

Collaborator **Chandrika Kamath,** CASC

A common trend in 2D object recognition is to detect and leverage the use of sparse, informative feature points. The use of such features makes the problem more manageable while providing increased robustness to noise and pose variation. In this work, we extend these ideas to the spatiotemporal case. For this purpose, we show that the direct 3D counterparts to commonly used 2D interest-point detectors are inadequate and propose an alternative. Anchoring off of these interest points, we devise a recognition algorithm based on spatio-temporally windowed data. We present promising recognition results on a challenging real-world database of mouse behaviors.

We attempt to extend the above approaches developed for object recognition to the problem of behavior recognition, i.e., from the spatial to the spatiotemporal domain. These extensions are not always direct, but rather follow the general spirit of using sparsely detected features for object recognition. We propose to characterize behavior through the use of spatiotemporal feature points.

For the above purpose, we show that the direct 3D counterparts to commonly used 2D interest-point detectors are inadequate and propose an alternative. We extend descriptors for spatial interest points to cuboids of spatiotemporally windowed data. Cuboids extracted from a number of sample behaviors are clustered and the resulting cluster centers serve as a dictionary of atomic units of behavior we call "actons." The only information kept from all subsequent video data is the location and type of the actons present. We argue that such a representation is sufficient for recognition and robust with respect to variations in the data. We show an application of this framework, utilizing a simple classification scheme, to a challenging real-world database of mouse behaviors.

Given these feature vectors and labelled training data, we trained two different classifiers: linear discriminant analysis and a DAG support vector machine using a radial basis kernel. Both were applied after dimensionality reduction using principal components analysis. As mentioned in each case,six clips were used for training and the seventh for testing. We limited the amount of training data for the over-represented categories, such as exploring, otherwise categories such as drinking would be drowned out. However, we used all the testing data. The SVM classifier outperformed LDA by a few percentage points. The classifiers achieved an accuracy ranging from about 75% to 95%.

Robust Trajectory and Appearance-Based Data Association for Multi-Object Tracking

Principal Investigator **Michael C. Burl,** University of Colorado, Boulder

Collaborator **Chandrika Kamath,** CASC

Reliable tracking is clearly a prerequisite for more advanced video-mining operations. This study builds upon our previous work in multi-object tracking from surveillance-style video sequences. The primary goal was to improve upon the tracker's ability to assign and maintain an object's unique track ID for the duration of time that the object appears in the scene.

The general structure of our previous tracker consisted of six steps.

1. Background estimation and subtraction
2. Thresholding and spatial grouping of detected pixels into blobs
3. Gating and data association to link blobs with tracks
4. Birth of new tracks if there are unassociated blobs
5. Measurement update (incorporating new observations into the state estimates for each of the tracks)
6. Time update (predicting the position and appearance of the tracked objects at the next observation time).

A key problem with this type of tracker is that the blobbing process is not perfect and leads to over-segmentation (a single object split into multiple blobs) and under-segmentation (multiple objects lumped into a single blob). Subsequent steps in our previous tracker could not properly deal with these fundamental grouping errors.

In this study, we investigated a number of ideas that we believed could improve tracking performance. A two-pronged approach was pursued with part of the effort devoted to getting more information from an object's trajectory and part of the effort devoted toward getting more information from an object's appearance.

The following component ideas were considered.

- Compensating for perspective effects
- Using more refined dynamical models
- Using decision-theoretic local search (structured change detection)
- Establishing precise correspondence between parts of a tracked object from one frame to the next, splitting and merging blobs
- Using particle filters to overcome some of the limitations caused by hard data association decisions.

Although we explored these ideas in some detail, we were not able to integrate these disparate components together into a unified tracker within the allotted time, nor were we able to conduct systematic evaluations over benchmark data sets to characterize the overall impact on performance. However, we plan to continue working toward these two goals.

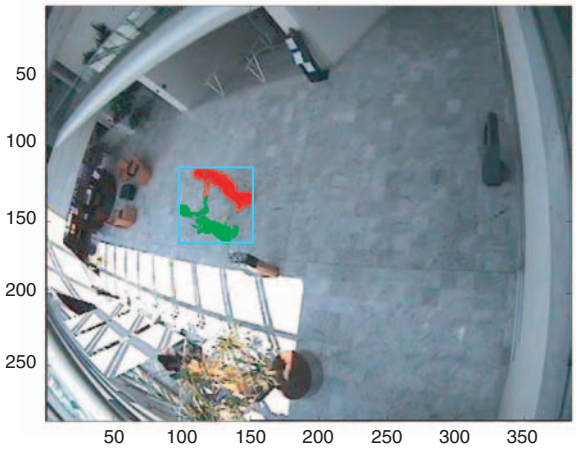

Two people meet and shake hands. The object detection component initially groups the two people together into a single blob (cyan bounding box), but the tracker is able to use the appearance and trajectory information from the previous frames to correctly split the blob into two individuals.

Multi-Constraint and Multi-Objective Partitioning-Complex Networks

Principal Investigator

Umit Catalyurek, Ohio State University

Collaborator

Edmond Chow, CASC

The aim of this project was to get better insight into partitioning properties of complex networks. With the advances in high-performance computing and high-resolution sensors in many fields of the computational sciences, data sizes are getting bigger and bigger. The sizes of complex networks that arise in homeland security applications are no different than any the other applications. Both the size and on-demand query and analysis requirements of applications that use complex networks necessitate both parallel and distributed computing. To achieve good performance, data needs to be partitioned in a load-balancing manner and communication requirements needs to be minimized.

Starting with random networks, we investigated the performance of 1D and 2D coarse-grain hypergraph partitioning of various networks provided to us by LLNL. Using the Performing Tools for Hypergraph (PaToH) partitioning tool, we analyzed the change in two cost metrics—graph edge cut and hypergraph connectivity-1—with respect to several graph parameters and the number of partitions. We also derived theoretical relationships between those parameters and the cost metrics.

In the experimentation phase for the type of the networks that have a generator, we generated multiple graphs with different parameter values, such as average degree and number of vertices. For each graph, we first used a random partitioner to generate a partition that would be used as a baseline to evaluate the performance of PaToH. Next, we partitioned these graphs using 1D and 2D coarse-grain hypergraph models using PaToH as our partitioner. Finally, we compared the performances of those partitioning schemes and explained the results in accordance with the theoretical predictions.

Our results showed that both 1D and 2D partitioning achieve much better results than the random partitioning, even though both approaches produce similar communication volume requirements. However, in applications that are sensitive to the number of communications, partitioning via 2D coarse-grain decomposition achieves much better results, providing better performance. Searching through paths in complex networks, especially for the homeland security applications that motivate our work, provides a practical real-world scenario where partitioning using 2D coarse-grain decomposition would be advantageous.

Numerical Simulation of Dispersion of Solid Particles Moving at Supersonic Speeds in Turbulent Flows

Principal Investigator

Said Elghobashi, University of California, Irvine

Collaborator

Bill Bateson, CASC

This work aims at developing a mathematical model that predicts the trajectories of spherical solid particles moving at supersonic speeds in a turbulent flow. The mathematical model involves the numerical solution of the three-dimensional, time-dependent, compressible, Navier–Stokes and energy equations. The project comprises two main tasks to be performed in parallel. Each task will require four years for completion.

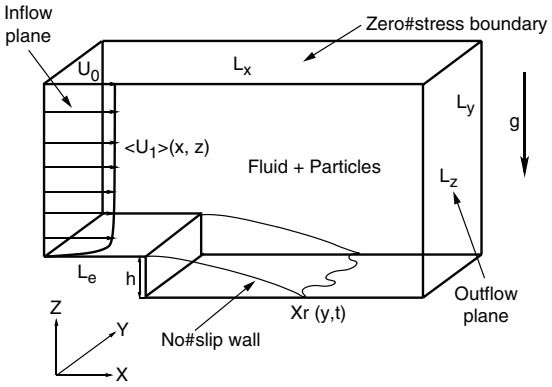

Figure 1. Schematic of the computational domain.

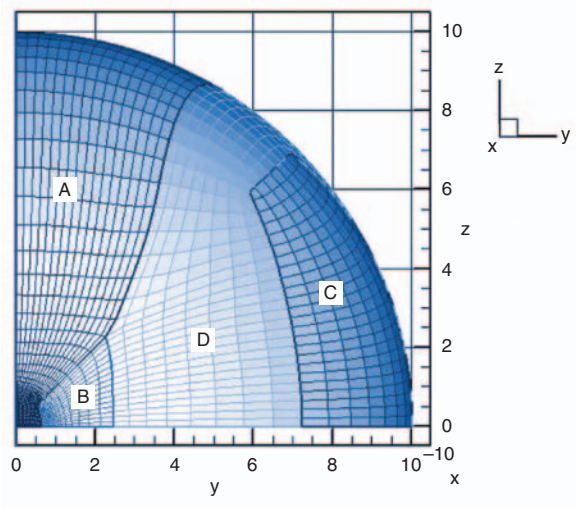

Figure 2. Grid surface ZY plane view

Task 1. The objective of this task is to compute the unsteady 3D compressible flow around a single-fixed sphere subjected to an air flow of a speed about 5000 m/s. An existing code at University of California, Irvine has been modified for a laminar compressible flow to account for the presence of a fixed sphere. Integrating the forces imparted on the sphere surface provides the values of the drag and lift on the sphere. The effects of varying the Mach number and particle Reynolds number will be investigated. A differential equation of motion of a sphere at supersonic speeds ($M > 1$) will be developed based on the computed drag and lift forces.

Task 2. The objective of this task is to compute the 3D turbulent isotropic compressible flow laden with many spherical solid particles moving at supersonic speeds in a cubical domain. The motion of each particle will be governed by the differential equation mentioned above in Task 1. A mathematical model for the collision between the particles will be included in the particle motion equation.

The detailed report describes the mathematical approach and the results of an incompressible flow over a single sphere. The computed values of the drag coefficient agree well with those of Kim, et.al. for the range of Reynolds numbers Re = 20 to 100.

Analysis and Visualization of Scientific Data Sets Using Generalized Segmentation Methods

Principal Investigator **Bernd Hamann,** University of California, Davis

Collaborator **Valerio Pascucci,** CASC

Segmentation of scientific data sets is becoming increasingly important in the context of understanding massive numerically simulated data sets at higher levels of qualitative behavior. Topology is one area in mathematics that allows us to characterize the behavior of scalar-valued data sets based on analyzing critical points and their relationships. This effort developed new techniques and program modules that support the automatic and mathematically sound segmentation of scalar-valued data sets into subsets (subregions of the physical domain) of topologically (and thus qualitatively) distinct behavior.

Traditional volume rendering of a data set involves visualizing surface properties or utilizing a transfer function. An alternative approach visualizes topological information, supporting a more qualitative understanding. The Morse–Smale complex is a structure that represents the topology of a scalar data set. The Morse–Smale complex can be simplified for use in applications, such as topological smoothing or hierarchical data visualization.

We have developed algorithms for extracting and simplifying Morse–Smale complexes that rely solely on combinatorial decisions, therefore avoiding numerical instability. We extract a Morse–Smale complex from a scalar field defined over a tetrahedral mesh using a multi-stage, region-growing approach. We also simplify a Morse–Smale complex through the cancellation of critical point (zero-gradient-point) pairs. There are two types of cancellations—saddle point–extrema cancellations and 1-saddle–2-saddle cancellations. We have developed rules for determining when a cancellation is valid and rules for reconnecting a simplified complex.

A multi-resolution representation of a Morse–Smale complex can be obtained through a simplification hierarchy. We have defined rules for valid simplifications of a complex and developed a hierarchy based on independence of cancellations. We declare two cancellations to be independent when their affected areas in the Morse–Smale complex do not overlap. We also represent a hierarchy as a directed acyclic graph that encodes the independence of cancellations. A multi-resolution reconstruction of the complex can be obtained by cutting the hierarchy graph.

Additional work has been done under this subcontract for error estimation in the context of rendering large multi-resolution data sets. We have developed an algorithm to estimate the error associated with changing resolution levels in a data set stored as a multi-resolution hierarchy. Our method utilizes a pre-processing step that calculates intermediate resolution levels based on error to allow for interactive modification of transfer functions used in a volume rendering application. Our approach balances storage overhead cost and quality of error estimates in support of an efficient method for choosing the subset of a volumetric data set that will be rendered in a near-optimal way.

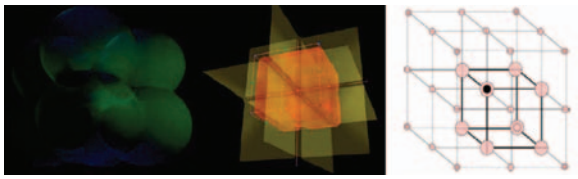

Iso-surface of test data set, stable and unstable manifolds, and associated Morse–Smale complex.

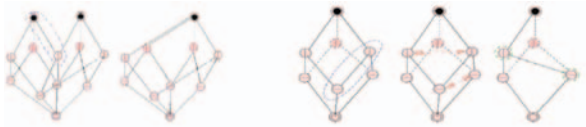

A saddle-maximum cancellation (left) and a 1-saddle–2-saddle cancellation.

Element Agglomeration AMGe for Contact Problems

Principal Investigators

Ana H. Iontcheva and Randolph E. Bank, University of California, San Diego

Collaborator

Panayot S. Vassilevski, CASC

The goal of this project was to develop algebraic multigrid methods for solving constrained-minimization problems, mainly contact problems in linear elasticity, discretized on general unstructured meshes, using the finite element method, in particular

- Signorini's problem
- Two-body contact problem
- Obstacle problem

For the solution of the Signorini's problem—contact of a linearly elastic body with a rigid frictionless foundation—we have developed two multilevel algorithms

- Multilevel subspace minimization algorithm
- FAS-constrained optimization algorithm

These schemes utilize element agglomeration coarsening away from the contact boundary, which allows for a straightforward construction of coarse-level approximations that automatically satisfy the fine-grid constraints. The two algorithms provide monotone reduction of the energy functional throughout the multilevel cycle. A code (in C++) has been developed and a paper has been written for the solution of the Signorini's problem. The results from this paper have been extended to 3D.

For the solution of the two-body contact problem, contact of two linearly elastic bodies in a mortar-method-based algorithm is the target of our future research and development. It will utilize the already developed methods for a single-body Signorini's contact problem.

Hypre for Symmetric Generalized Eigenvalue Problems

Principal Investigator

Andrew Knyazev, University of Colorado, Denver

Collaborator

Charles Tong, CASC

Symmetric eigenvalue problems are crucially important in structured mechanics and electronic structure calculations. The goal of the project is to implement a modern eigenvalue solver into Hypre in order to take advantage of Hypre high-quality preconditioners for parallel clusters with a user-friendly interface. The Locally Optimal Block Preconditioned Conjugate Gradient (LOBPCG) method for symmetric eigenvalue problems has been chosen for implementation as one of the most promising eigensolvers in the current literature. The LOBPCG apparently combines the fast convergence of the block's Lanczos method with the simplicity and robustness of the block's steepest descent and allows the preconditioner to be plugged directly into the eigensolver without any inner-outer iteration schemes. The result is a well-written native Hypre C implementation of LOBPCG that works in three Hypre interfaces and is almost as efficient as the Hypre PCG linear solver.

The basic functionality of the code has been tested on University of Colorado, Denver and LLNL clusters on a number of test problems, including non-self-consistent electronic structure calculations eigenproblems, and the code is checked into the Hypre Alpha. Presently, this is the first and only preconditioned eigenvalue solver for parallel clusters.

An entirely new implementation of the LOBPCG for eigenvalue problems has been written as a part of the Hypre, consistent with Hypre Krylov-based linear solvers. The distinctive features of the new LOBPCG code are as follows.

1) The code is rewritten from scratch in order to implement the **lopbcg.c** code at the same level of abstraction as the Hypre **pcg.c** code.
2) The LOBPCG now works for three Hypre interfaces—IJ, Struct and SStruct—and includes test drivers for every interface. As of August 31, 2004, the LOBPCG drivers reflect the latest changes in the corresponding Hypre drivers and are prepared to serve as combined drivers for both linear systems and eigenvalue problems.
3) A new Hypre type, hypre_MultiVector, is introduced, and a preliminary implementation of Hypre MultiVector functions based on the existing parallel vector types (hypre_ParVector, hypre_StructVector and hypre_SStructVector) is used in the LOBPCG code. A significant future effort by the Hypre team is necessary to turn it into an actual MultiVector with efficient implementation of Hypre MultiVector functions by eliminating redundant MPI calls and by using Basic Linear Algebra Subprograms (BLAS).
4) The code solves generalized eigenvalue problems, as well as regular eigenvalue problems.
5) The code allows the use of constraints.

Preconditioning of Finite-Element Saddle-Point Problems

Principal Investigators

Raytcho D. Lazarov and Joseph E. Pasciak, Texas A&M University

Collaborator

Panayot Vassilevski, CASC

The goal of this research is the efficient solution of the discrete equations that result from finite-element or finite-difference approximation of partial-differential equations of mathematical physics. This involves the development and analysis of algorithms especially tailored for execution on medium- to large-scale parallel computing platforms. Our approach involves the application of theoretical tools from the analysis of partial differential equations to motivate and analyze new computational algorithms.

For implementation of highly accurate methods and efficient algorithms, we need approximation methods that provide greater flexibility in the grid generation process, increase the portability of various approximation methods and computer implementations, enhance the capabilities of coarsening strategy in parallel algebraic multi-grid methods, and provide a natural and practical way for parallel domain decomposition methods and parallel adaptive methods based on a posteriori error analysis. Our investigations produce competitive algorithms that can be used in various codes for complex applications in physics and engineering.

Main results of the research

1. A new, inexact Newton algorithm for the solution of second-order problems with higher-order nonlinearities was proposed and analyzed. This approach was based on a stability analysis in Sobolev spaces of order greater than one.
2. A new multigrid algorithm was proposed and analyzed for an electromagnetic problem. This method involves strengthening a curl-curl term by adding a discrete "grad-div" term. The resulting two-level multigrid algorithm was analyzed.
3. A stabilization framework for the discontinuous Galerkin finite-element method was developed. It was implemented for Raviart–Crouzeix non-conforming finite elements and tested on three-dimensional problems of elasticity.
4. A new stable scheme of exponential fitting type was proposed and studied for convection-diffusion problems.

Biological Fluid Flow in Micro-Electro-Mechanical Systems

Principal Investigator

Dorian Liepmann, University of California, Berkeley

Collaborator

David Trebotich, CASC

Our investigation aims to characterize the flow behavior of biological macromolecule solutions in silicon microfluidic devices. Digital Particle Image Velocimetry (DPIV) is used to quantify the velocity fields of DNA-laden solution flow under microfluidic conditions. Deviations from Newtonian flow fields brought about by the viscoelastic fluid rheology, concentration effects and conformational changes of the molecules will be assessed for flows through a variety of microfluidic geometries, such as straight channels, contractions, and expansions. Experimental results will be used to validate a computational design tool for bio-detection microdevices being developed at LLNL.

The characterization of flows containing high concentrations of macromolecules is critical for the optimal design of microfluidic systems for biochemical analyses. Since the fluid flow behavior at the microscale differs from the macroscale, a simple scaling down of processes may be insufficient to elicit the same performance. This thereby motivates research directed toward the characterization of the microflow behavior of the biological macromolecule DNA, commonly used in biochemical assays. Studies have shown that λ-DNA can undergo dramatic stretching in microfluidic flow with conformational changes occurring in both elongational and shear flows (Shrewsbury et al. 2001). The device geometry and flow rate, as well as the viscosity, concentration and relaxation time of the solution, influence the conformation of the DNA.

A 1.67:1 microfabricated abrupt planar contraction design was chosen as a test device (see figure). Experimental pressure drops for 450 μg/ml DNA (semi-dilute) solutions across abrupt contraction were measured. Velocity fields for flows of Newtonian and semi-dilute DNA solutions through abrupt contraction and abrupt expansion (inverted contraction device) were quantified using DPIV. No flow regions were observed near the contraction corners at slower flow rates and there is visual evidence of recirculation at faster flow rates.

Validation of the pressure testing system was conducted by measuring water flow through a narrow tube and comparing to theoretical solution. Pressure drop data was collected for flows of Newtonian (water) and non-Newtonian (DNA-laden) solutions through gradual and abrupt contraction devices. The effects of DNA solution concentration were explored and the onset of elastic, non-Newtonian behavior determined. DPIV was used to quantify flow fields for water flow in straight rectangular channels. A comparison study was conducted to determine the minimum interrogation region size for accurate data processing.

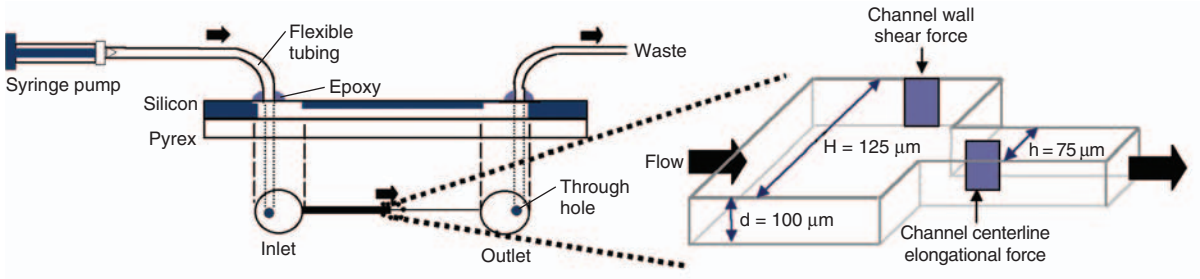

A microfabricated abrupt planar contraction design test device.at a ratio of 1.67 : 1.

TOC

Enabling Large-Scale Data Access

Principal Investigator

Ling Liu, Georgia Institute of Technology

Collaborator

Terence Critchlow, CASC

One of the ultimate goals of the Enabling Large Scale Data Access project is to produce a fully automated, end-to-end wrapper code generator through the design, development, and integration of service class descriptions with the XWRAP systems. The main idea is to provide mechanisms that enable XWRAP systems to take a generic description of a class of search interfaces (service classes), as well as the URL of a particular interface that is a member of the given class, to produce a functional wrapper that will take a class-specific query and produce XML views of the query results obtained through this particular interface. In the second year of this project, we focused on three main efforts.

(1) Extend the Web Spider developed under the previous subcontract to generate a set of Java 1.4 wrappers for interfaces that it identifies as matching a service-class description (SCD). Collectively, these wrappers will query all valid combinations of user input as specified by the SCD and the parameters enumerated by the interface. Each wrapper will be able to handle multiple types of errors (e.g., timeouts, connection refused, page not found) and return an appropriate message to the calling routine. The new version of the Web Spider is called DynaBot.

(2) In the DynaBot development, we have updated the Web-crawling component of the spider to be consistent with the interface identification component, updated the httpUnit library, and resolved JavaScript errors that prevent it from being used against some popular bio-portal sites, such as the National Center for Biotechnology Information's Basic Local Alignment Search Tool (BLAST) interface.

(3) We have tested the spider's ability to interact with the interface identification component by crawling the Web, starting at the dbCAT Web page and comparing the interfaces it identifies and the wrappers it generates with those identified and generated manually. This includes identifying and removing all obvious problems with the previous implementation of the spider, especially those identified in the second-year effort.

DYNABot System Architecture

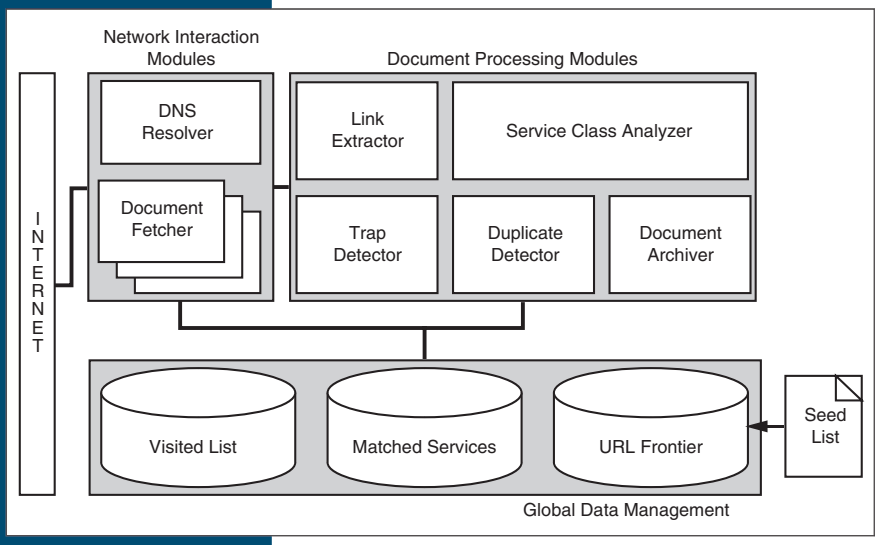

The components that make up a crawler are divided among three major component groups: network interaction modules, global storage and associated data managers, and document processing modules. The simplest crawlers require mechanisms for retrieving documents and determining if a particular URL has been seen. More advanced crawlers will include features like mirror-site detection and trap-avoidance algorithms. DynaBot utilizes an advanced crawler architecture for source discovery and adds a document processor that can determine if a dynamic Web source is related to a particular domain of interest. The architecture of Dynabot is shown at left.

For a more detailed description of this research, see the full report at: **http://disl.cc.gatech.edu/LDRD/**

Large Graph Visualization

Principal Investigator

Kwan-Liu Ma, University of California, Davis

Collaborator

Marvin Christensen, NAI

The Computer Incident Advisory Capability (CIAC) currently collects a great deal of data from sensors placed on a handful of computer networks. These sensors record activity on the networks so that the recorded information can be analyzed to provide useful information about both normal and anomalous activity on the networks.

This project focuses on one of the key significant technical challenges for interpreting this data. That challenge is a task in data mining—given the huge pool of data provided by the sensors, how can the interesting features be identified and classified? We address the problem by constructing a visualization system implementing novel methods of data analysis and information display. The process of exploring the space of a large data set using visualization is known as "visual data mining." It takes advantage of the pattern recognition facilities of the human visual system to detect patterns and anomalies in visual representations of abstract data.

We have developed a prototype system for visualizing TCP port data. The system permits analysts to discover the presence of any network security event that causes significant changes in the activity on ports. Since we currently only have access to very high-level data, the system is a very high-level tool and is useful mostly for uncovering high-level security events. Security events that consist of small details—an intrusion that includes only a few connections, for instance—are unlikely to be caught using this prototype system. Furthermore, since we have only obtained counts of activities rather than records of the activities themselves, the analysis can only go so far. The system can help identify suspicious traffic patterns, but it cannot see the traffic that caused the patterns. This is still useful, however, because analysts using this system can send the suspicious traffic signatures to analysts that have access to the full set of network traffic logs. The figure displays the current system interface.

Even in settings where only generalized information is available concerning network activity, we found many types of malicious activity can still be discovered using visualization. We have developed a tool that takes general, summarized network data and presents multiple, meaningful perspectives of the data and have demonstrated that this visualization leads to useful insights concerning network activity. Port scans of several types have been successfully detected, and many suspicious traffic patterns on individual ports have been uncovered. In addition, useful information about overall network traffic has been revealed, such as the rhythm of the traffic on commonly used ports as time progresses and the relationships between the various metrics used to describe port activity.

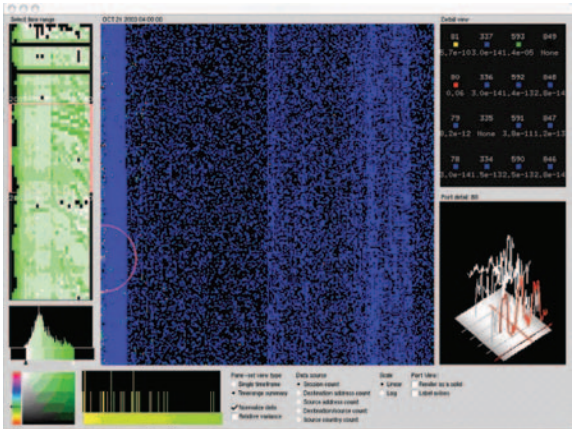

User interface of the TCP port data visualization system. The interface supports a drill-down process by allowing the analyst to explore high-level to lower-level information, from left to right, respectively.

Multi-Resolution Interactive Rendering of Large Scientific Data Sets Using an Image Cache

Principal Investigator

Kwan-Liu Ma, University of California, Davis

Collaborator

Valerio Pascucci, CASC

Large-scale simulations running on parallel supercomputers can produce very high-resolution information of the model phenomena, giving scientists the power to examine the phenomena at greater fidelity. Visualization enables scientists to better understand large amounts of data from such simulations. The problem with visualization is that conventional methods fail to handle the large amount of data. The data can take up too much system memory, as well as a considerable amount of time to render. Scientists are in need of tools that interactively search for and discover regions of interest within their simulations. Waiting for a long batch-mode visualization is not an option. In this project, we aimed at developing an interactive visualization solution enabling scientists to freely and effectively explore their simulation data.

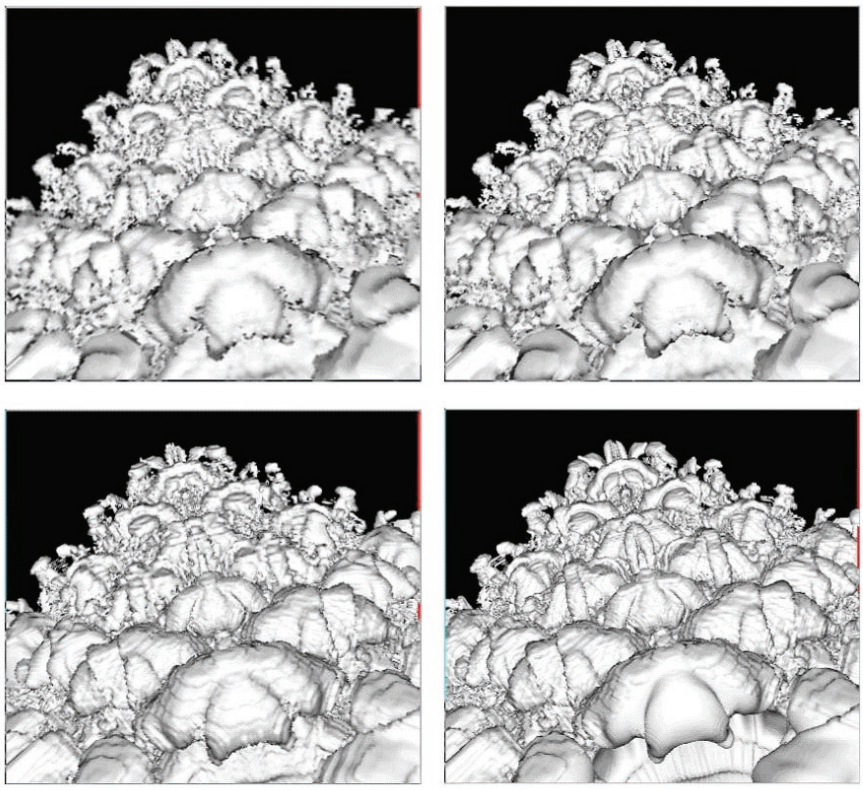

An example of the same scene rendered at different data resolutions. It is clear to see that as data resolution increases, images quality is greatly improved.

This project began as an attempt to provide a system that is capable of rendering large multiresolution data sets while taking advantage of a cluster computer. Our work involved the extension of a multi-layered image cache system (MLIC) developed at LLNL in several ways. First, the rendering engine was modified to access data sets stored in the Visualizations Schemes for Ultimate Scalability (ViSUS) cache oblivious to multi-resolution representation. This has improved the performance of the rendering engine with minimal loss in accuracy. Second, we removed the reliance on the Visualization Toolkit (VTK), which has greatly simplified the system implementation and its performance. Next, the system now runs on LLNL's new cluster computers, making it possible to use high-resolution output devices like a large PowerWall display. Finally, the user interface was revamped to be more intuitive and easy to use. The new system gives an easy method of navigating through a large dataset while simultaneously creating high-quality visualizations.

The extended MLIC system is a capable multi-resolution rendering system. It not only retains but enhances system interactivity, which allows a user to navigate through a dataset while it is being rendered. Rendering can be stopped and settings can be modified before completion, allowing a scientist to make changes on the fly. The new rendering classes have made the system render independent, which allows for the use of almost any rendering algorithm or rendering software.

As seen in the figure, the system can render a very coarse-resolution image, as well as a clear high-resolution image and several levels in between for the large Richtmyer-Meshkov turbulence simulation. This allows a scientist to have an idea about what is going to be rendered and possibly make changes. This is made possible by decoupling the rendering and displaying of images, allowing the system to always remain responsive even as the rendering is taking place. Responsiveness leads to easy navigation.

Open Source Software Technology for Transforming Scientific Problems

Principal Investigator **John Mellor-Crummey,** Rice University

Collaborator **Dan Quinlan,** CASC

Two leading software systems for source-to-source transformation and optimization of scientific programs are the ROSE infrastructure being developed at LLNL and the Open64/SL infrastructure being developed at Rice University. These efforts are complementary. The Open64/SL infrastructure has focused on software support for parsing and transforming Fortran 90-based programming models. ROSE has focused on software support for parsing, analysis and transformations of C++ based programming models. This project will build an open-source software technology to bridge these two infrastructures to allow ROSE's transformation capabilities to be applied to Fortran 90-based programs.

The project began with a study of Open64's WHIRL infrastructure, a study of LLNL's ROSE infrastructure (including its SAGE IR) and the construction of test programs to familiarize ourselves with ROSE. Currently, we have begun constructing a framework for the translation between WHIRL and ROSE. We have begun adding a C++ namespace to the ROSE code so that it can interoperate with the Open64 infrastructure and WHIRL. The translator itself is being constructed by modifying code that performs a WHIRL-to-WHIRL transformation to compile Co-array Fortran programs.

A Perl script will orchestrate the WHIRL-to-ROSE translation. This script will invoke mfef90 (the Open64 front end for Fortran 90) on a Fortran 90 source file to generate a file that contains the WHIRL representation source. The WHIRL external representation will be read by the translator proper (under construction), which will walk the WHIRL abstract syntax tree and translate each WHIRL operator into its corresponding operator in the SAGE IR. Once in the SAGE IR, ROSE-based tools can be applied to transform the resulting code. Next, the process will be reversed. The SAGE internal representation will be traversed and WHIRL will be reconstructed. The reconstructed WHIRL will be written out in a file using its external representation. Finally, the driver script will invoke Open64's whirl2ftool to regenerate Fortran from WHIRL.

We have been using Los Alamos National Laboratory's Parallel Ocean Program (POP), the National Academy of Sciences (NAS) parallel benchmarks, and the Adjoint MIT Ocean General Circulation Model to exercise Open64's support for analysis and transformation of Fortran. Currently, we are working with Open64's infrastructure on the parsing and regeneration of the POP code. With the exception of a minor problem unparsing arrays of structures, Open64 regenerates legal Fortran 90 code for POP.

Over the last several months, we have resolved a half dozen issues in Open64's support for Fortran, including correcting the very-high level WHIRL intermediate representation produced by the Fortran 90 parser, correcting unparsing of formatted I/O, correcting unparsing of array initializations. Currently, work is under way to correct problems with the constant pool produced in the WHIRL intermediate form and to produce a general infrastructure for correctly unparsing hierarchies of structures and arrays of structures.

Once we resolve known correctness issues, we will begin assembling a suite of codes for regression testing of Open64.

A Tightly Coupled Particle–Fluid Model for DNA-Laden Flows in Complex Microscale Geometries

Principal Investigators

Gregory H. Miller, University of California, Davis
P. Colella, D. T. Graves, D. F. Martin, P. O. Schwartz, LBNL

Collaborator

David P. Trebotich, CASC

We present a stable and convergent method for the computation of flows of DNA-laden fluids in microchannels with complex geometry. The numerical strategy combines a ball-rod model representation for polymers tightly coupled with a projection method for incompressible viscous flow. We use Cartesian-grid embedded boundary methods to discretize the fluid equations in the presence of complex domain boundaries. A sample calculation is presented showing flow through a packed array microchannel in 2D.

Approximate DNA in an actual microfluidic device is used for extraction. The left boundary condition is plug flow with a velocity of 0.1 cm/s; the right boundary is outflow (homogeneous Neumann); the top and bottom boundaries and the interior circular boundaries are a solid wall. The polymer is a 26-node approximation of DNA introduced near the left boundary as an inclined linear array after the fluid-flow field reached steady state. The polymer's trajectory causes it to become wrapped around the first circular element, where it remains pinned until the stochastic perturbations work it loose.

The fluid dynamic steps of this method are subject to an advective Courant–Friedrichs–Lewy stability condition only. The particle steps, without constraints, are also stable with this CFL timestep. When particles move far from the constraint manifold, however, the Lagrange multiplier algorithm of Ciccotti, et al. may diverge. We have found that the maximum particle displacement per time step for which the Ciccotti, et al. algorithm is stable may be extended for most systems by centering the constraint force at the conclusion of the timestep instead of the beginning.

We use a backward Euler time-stepping strategy that is formally first order accurate. To make the overall method second-order, it will be necessary to replace at least this with a Runge–Kutta time-stepping strategy.

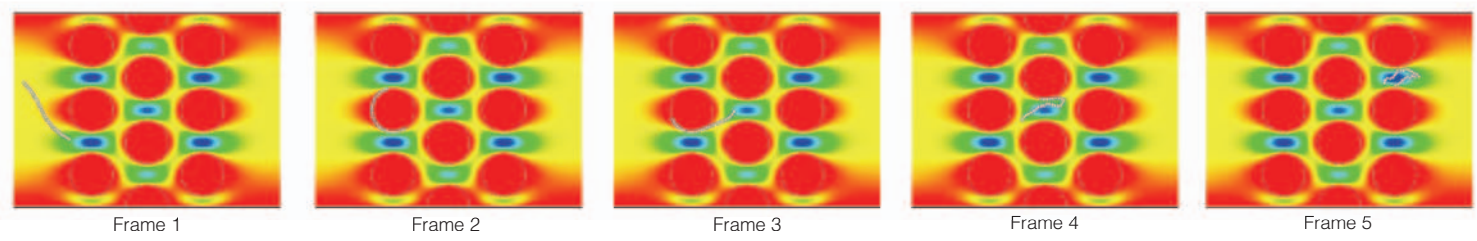

The time sequence of genomic DNA flowing in a 2D model of a packed-bed reactor PCR chamber. The DNA molecule enters from the left (frame 1), then wraps around a bead (frame 2), is loosened by hydrodynamic and Brownian forces (frame 3) and is swept out of the chamber by the flow field (frames 4 and 5). The color map indicates underlying flow field.

TOC

Clustered and Hierarchical Networks

Principal Investigator

Alex Pothen, Old Dominion University

Collaborator

Edmond Chow, CASC

We surveyed clustered and hierarchical representations of large networks available in the literature and made progress on two aspects of clustered graphs during the period of this contract.

First, we designed a new algorithm for discovering clusters in large networks that pays attention to the **local density** of the network. The density of a subgraph is the ratio of the number of edges joining nodes in the network to the maximum number of such edges possible (i.e., a complete subgraph or clique). In a large network, some subnetworks have high density, while others have low density, and a clustering algorithm that takes into account the local density would be able to do a better job of clustering.

In order to describe our algorithm, we need the concept of a ***k*-core**, a maximal subgraph in which every vertex is incident on at least k edges that belong to the subgraph (i.e., the degree of each vertex in the subgraph is at least k, when only edges in the subgraph are considered). The core number of a vertex is the maximum value of k as it belongs to a k-core, and the corresponding core is the maximum core to which the vertex belongs. The new clustering algorithm weights each vertex with the density of the connected component of the maximum core to which it belongs. The clustering proceeds by choosing a vertex with the highest weight to be the "seed vertex" for a cluster and then successively adds vertices whose weights are not too different from it to the cluster. After one cluster is complete, we begin with a new cluster in the same fashion.

This clustering algorithm has two advantages over existing ones. First, the weight of a vertex depends not on its degree but on the density of its neighborhood, thus clustering low-degree vertices together with high-degree vertices to which they are adjacent. Second, the algorithm can be efficiently implemented in time $O(K\,|E|)$, where K is the maximum core number of any vertex, and $|E|$ is the number of edges, in the network. We have implemented the weight computation phase of the algorithm and are currently working on the clustering phase.

Second, we are implementing a clustering algorithm for bipartite networks in which vertices in both parts need to be clustered. The algorithm relies on a singular vector of the vertex edge incidence matrix of the network. We followed the suggestion of Ravi Kannan, Santosh Vemula, and colleagues to speed up the singular-vector computation by computing with a submatrix of the incidence matrix. The columns in the submatrix are chosen with a probability proportional to the column weight. Then, the computed singular vector can be shown to be close to the singular vector of the larger matrix with high probability. Our preliminary results on the original incidence matrix shows that this approach works well, especially when combined with a combinatorial "local-smoothing" algorithm.

Spectral AMGe, ALE3D & FOSPACK

Principal Investigator

John W. Ruge, Front Range Scientific Computations, Boulder, CO

Collaborator

Rob Falgout, CASC

The majority of the work performed in FY 2004 has been on FOSPACK, a package for the discretization and solution of First-Order System Least-Squares (FOSLS) formulations of a range of PDEs. Algebraic Multigrid (AMG) is incorporated as the linear equation solver, and some of the effort has also gone into extending the range and improving the efficiency of that method. Progress has been made in several of the proposed areas of study and is summarized below.

AMG for Constraints

Slide-surface constraints were added to a two-dimensional, two-body elasticity problem. The constraints specified that two faces in contact remained in contact under application of a force, coupling the normal displacements along the two faces (and allowing for slippage in the tangential direction). Tests were conducted incorporating the constraints either strongly by essentially eliminating rows/columns from the matrix, or weakly by including them as weighted penalty terms in the functional to be minimized. Both resulted in an asymptotic V-cycle convergence factors that were nearly identical to those obtained without the constraints (where the two bodies were fully decoupled).

Parallel FOSPACK

Initial parallelization of 3D FOSPACK and coupling to Hypre/BoomerAMG has been completed for linear PDE systems with weak boundary conditions. The code has been ported to a 128-processor Beowulf cluster, and improvements in implementation resulted in better parallel performance. Treatment of nonlinear problems and coupled nonlinear multi-physics systems has also been improved. Tests performed on a coupled nonlinear fluid-elastic model and Maxwell's equations problem have resulted in solutions consistent with those obtained from the scalar FOSPACK code. The major effort for 3D FOSPACK is the introduction of higher-order discretizations, which are necessary for obtaining reasonable accuracy for 3D problems while keeping storage requirements within practical limits.

AMG Algorithm Enhancements

High-operator complexity can be a problem for AMG, especially in large 3D applications. In such cases, the coarse grid operators can be too large or too dense. Aggressive coarsening and multi-pass interpolation have been implemented in both the serial and parallel codes, resulting in markedly lower complexity, although some degradation of convergence is noted. A number of alternatives are being tested in the scalar code in an effort to both reduce complexity while retaining (or even improving) convergence.

Several variations of the basic AMG algorithm are being studied in an effort to improve its robustness. The use of compatible relaxation (CR) for choosing the coarser grids, combined with interpolation designed to minimize the trace of the coarse grid operator, looks quite promising. For some problems, the standard AMG interpolation is not accurate enough, particularly those with near-null space components. For this, an adaptive approach, in which interpolation is explicitly modified to better match the smoothest (slowest-to-converge) components, is being tested. The current approach is a hybrid of the recently developed adaptive Smoothed Aggregation (aSA) and a more straightforward adaptive AMG (aAMG).

TOC

Development of Efficient and Robust Algorithms for the Numerical Solution of Partial Differential Equations for High-Speed Reactive and Nonreactive Flow

Principal Investigator **Donald W. Schwendeman,** Rensselaer Polytechnic Institute

Collaborator **William Henshaw,** CASC

The work carried out under this subcontract involved the development and use of an adaptive numerical method for the accurate calculation of high-speed reactive and nonreactive flows on overlapping grids. In the reactive case, the flow is modeled by the reactive Euler equations with various choices for the equation of state and reaction rate model.

A numerical method has been developed to solve the nonlinear hyperbolic partial differential equations in the model. The method uses an unsplit, shock-capturing scheme, a Godunov-type scheme to compute fluxes, and a Runge–Kutta error control scheme to compute the source term modeling the chemical reactions. Two approximate Riemann solvers are now available and can be used for a general (non-ideal) equation of state. An exact Riemann solver is also available for an ideal, gamma-law equation of state.

Adaptive mesh refinement (AMR) has been implemented in order to locally increase grid resolution. The numerical method may also be applied to nonreactive flow problems, in which case the reactive source term is simply set to zero. The numerical method uses composite overlapping grids to handle complex flow geometries in either two or three dimensions. The code is part of the Overture–OverBlown framework of object-oriented codes, and the development has occurred in close collaboration with Bill Henshaw and Lori Diachin, and other members of the Overture group within the Center for Applied Scientific Computing.

During the period of this subcontract, a number of tasks were accomplished, including

- Further development of the numerical method for 2D moving meshes
- Implementation of additional Riemann solvers, including a Harten–Lax–van Leer (HLL) approximate Riemann solver for a general equation of state and an exact Riemann solver for an ideal equation of state
- An extension of the numerical method to handle axisymmetric flow
- An extension of the numerical method to handle multi-material, non-reactive flow

Workshops and Conferences

The ISCR-sponsored or co-sponsored 10 scientific workshops in FY 2004. Two of these were hosted locally and exclusively by the Laboratory. The rest were hosted in cooperation with other organizations, such as the Society for Industrial and Applied Mathematics (SIAM), the Institute for Pure and Applied Mathematics (IPAM), the Department of Homeland Security, or Argonne National Laboratory and held off site. Some ISCR workshops are one-of-a-kind exploratory workshops that assemble experts to scope out possible new programs. Others have become part of the fabric of their disciplines and are held at regular intervals. In each case, there is a vital LLNL interest and typically, several Laboratory researchers participate.

TOC

The 2003 International Conference on Preconditioning Techniques for Large Sparse Matrix Problems in Scientific and Industrial Applications

Dates October 27–29, 2003

Location Napa, California

The 2003 International Conference on Preconditioning Techniques for Large Sparse Matrix Problems in Scientific and Industrial Applications (Preconditioning 2003) was the third conference of its kind to focus on preconditioning techniques for solving sparse matrix problems. The first conference took place in Minneapolis, MN, in June 1999 and the second in Tahoe City, CA, at the end of April 2001. The first conference drew close to 100 participants and the second about 70 participants, while the third drew about 80 participants.

One of the characteristic themes of the meeting is its emphasis on real-life ("industrial") problems. For this reason, all three conferences enjoyed a healthy balance between academia and industry/government labs in its mix of participants. In addition, there is a rather important contingent of participants from overseas (mainly Europe). The preconditioning meetings have been quite successful and are now being viewed by the community as the premier specialized conference on preconditioners.

The Napa conference featured 7 invited speakers, 29 contributed papers and 8 posters. One of the goals of the meeting is to foster dialogue between practitioners and academics. As a rule, the organizing committee gives a charge to the program committee to nominate invited speakers with a goal of keeping a good balance between the number of talks on "methods" and those on "applications." The organizing committee then finalizes the selection to reach this goal. This particular meeting reached a good balance. There were four invited talks on applications and three invited presentations on algorithmic aspects.

One feature that distinguished this meeting from the previous two was that we allowed more parallel sessions. Previously, few (Tahoe City) or zero (Minneapolis) parallel sessions were scheduled. This change was necessary to increase the number of presentations as there were many high-quality abstracts.

There was a good mix of attendees from academia, research laboratories, and industries. In particular, the DOE was well represented; with more than 20 participants from all of the major DOE laboratories (ANL, LBNL, LLNL, LANL, and SNL). Industry participation was also relatively strong. Overseas participants came from Belgium, China, Denmark, France, Germany, Great Britain, Israel, Italy, Japan, The Netherlands, Russia, and Tanzania.

Multi-Algorithm Methods for Multiscale Simulations Workshop

Dates January 14–16, 2004

Location Livermore, California

The Multi-algorithm Methods for Multiscale Simulations Workshop was held January 14–16, 2004, at the Hilton Garden Inn in Livermore, California. The event was hosted by CASC, ISCR and LLNL. It was sponsored by the Institute for Terascale Simulation (ITS) on behalf of the ASC Program. In all, 45 attendees, including 26 from U. S. Department of Energy laboratories and the balance from academia and industry, participated in the two-and-one-half day workshop.

Multiscale simulation is a central emerging numerical modeling paradigm for many science and engineering problem areas at LLNL and throughout the scientific community. Application areas include: materials design (nano-wires in computer chips, photonics, micro-electromechanical systems); biological systems (protein docking); medicine (drug delivery systems); and nuclear and aerospace technologies where materials failure and response in severe environments is a primary concern (dislocation patterns in fatigue and creep, surface roughening and crack nucleation in fatigue, etc.). Important scales in such problems range from macroscopic, where continuum models based on differential equations are usually employed, to atomistic, with quantum mechanical models applied at the finest resolution, plus all intermediate scales.

The Multi-algorithm Methods for Multiscale Simulations Workshop comprised five half-day sessions on various multiscale topics. Sessions were sorted by the organizers into methods and applications in solids, liquids, and gases and cross-cutting techniques. Each session concluded with a panel discussion in which the speakers and audience interacted richly and informally.

An important aim of the workshop organizers was to gather practitioners of multi-algorithm methods from various fields to find common ground in their work and engage in intellectual cross-fertilization. For example, materials scientists and gas dynamicists are not accustomed to working closely together, let alone with the mathematicians and computer scientists also present, since traditional conferences and meetings are usually focused on specific disciplines or problem areas. Many of the workshop participants explicitly stated that they found the workshop format and technical exchanges refreshing and informative. Several participants expressed strong interest in making the workshop an annual event.

Common challenges identified at the workshop included locating relevant scale boundaries and in grafting together representations of the solution from different methods in such boundary regions in which both are valid, in order to model a global system for which no single method is everywhere valid or efficient. Participants also discussed important challenges in understanding physical and mathematical error analysis in hybrid computational models and implementation challenges associated with complex models for large-scale parallel platforms.

An especially important outcome of the workshop is that DOE and academic researchers gained familiarity with many commonalities, as well as differences, among the problems on which they work. For example, it was revealed that understanding dissipative processes in solid mechanics (e.g., dislocations) can benefit from work done in modeling liquids. Also, the principle of entropy production used to understand the breakdown of continuum models in liquid and gas simulations appears to possess similarities with the principle of power dissipation used in solid mechanics. Other important areas identified as requiring further exploration included methods for bridging widely disparate time scales in hybrid simulations and methods for fluid-structure/gas-surface interactions, as well as incorporating more detailed chemistry in biological and nanoscale problems.

https://www.llnl.gov/casc/workshops/multiscale_simulations

Copper Mountain Conference on Iterative Methods

Dates March 28 – April 2, 2004

Location Copper Mountain, Colorado

The Copper Mountain Conference Series, held annually in early April at Copper Mountain, CO, alternates subjects between Multigrid Methods in odd-numbered years and Iterative Methods in even-numbered years. It represents an important forum for the exchange of ideas in these two closely related fields.

The Copper Mountain Conference on Iterative Methods was held March 28 – April 2, 2004. A total of 185 mathematicians from all over the world attended the meeting, which began with a reception on Sunday, March 28. During the following 5 days of the meeting, 131 talks on current research topics were given. Talks were organized into the following sessions.

1. Multigrid Solvers and Algebraic Multigrid
2. Saddle-Point Solvers
3. PDE Methods
4. Preconditioning Methods
5. Eigenvalue Methods
6. Multi-Physics Solution Methods
7. Krylov Subspace Methods
8. First Order System Least Squares Methods (FOSLS)
9. Nonlinear Solvers
10. Continuation Methods
11. Stochastic Systems
12. Parallel Algorithms
13. Applications
14. Software

In addition to the regular sessions, three evening workshops were offered. Monday night's workshop was a mini-symposium organized by Henry Tufo, who represented NCAR and the University of Colorado. Tuesday night, Michael Heroux from SNL organized a workshop on Sandia's Trilinos project. Wednesday night, a workshop organized by Eldad Haber from Emory University highlighted PDE-Constrained Optimization.

The sessions were all very well attended. The Copper Mountain Conference Series is known for having a very relaxed atmosphere and for fostering open, active discussions. This collaborative environment has characterized the meeting since the series began in 1983, and is one of the reasons many attendees come back repeatedly.

A student paper competition was held to stimulate student participation in the Conference. A panel of judges made up of members of the Program Committee selected the winners: Yair Koren (Technion, Israel); Ruth Holland (Oxford University, England); and Andrei Draganescu (University of Chicago).

http://amath.colorado.edu/faculty/copper/2004

Statistics and Practical Applications of Data Mining: Highlights from SDM04

Dates April 22–24, 2004

Location Orlando, Florida

The Fourth SIAM International Conference on Data Mining, held in Orlando, FL, April 22–24, 2004, continued the tradition of providing an open forum for the presentation and discussion of innovative algorithms, as well as novel applications of data mining. A record number of paper submissions this year marked not only a growing interest in the field, but also a greater acceptance of the conference among data mining researchers and practitioners.

Student authors accounted for a large percentage of the accepted papers, and their papers were reviewed under the same stringent guidelines as regular papers. The best student paper award was given to Martin Law from Michigan State University for his work on manifold learning. The award for the best algorithms paper went to a team from the University of Texas at Austin for their work on clustering, while the best applications paper was on enhancing communities of interest by a team from AT&T Laboratories.

A running theme of the conference was the practical application of data mining, including opportunities in various problem domains and practical lessons learned by those solving real data analysis problems in these domains. This was reflected in the topics covered in the three tutorials: analysis of patients' medical data, data mining for computer security, and mistakes commonly made in data mining and ways to avoid them.

In an industry–government session, speakers discussed problems encountered in the telecommunications industry, the role of information visualization, and data mining in such diverse domains as aviation safety and security, performance of computer networks, and earth sciences. Applications of data mining were also the subject of three of the keynote talks: Sara Graves of the University of Alabama at Huntsville considered issues of data usability; David Page of the University of Wisconsin Medical School elaborated on data mining questions raised by biology data; Ted Senator from the Defense Advanced Research Projects Agency (DARPA) discussed "connecting the dots." The increasing importance of homeland security was also reflected in many of the conference workshop topics, which ranged from link analysis, counterterrorism and privacy, to data mining in resource-constrained environments. More traditional topics, such as bio-informatics, mining of scientific and engineering data sets, and high-performance and distributed mining, also continued to attract participants.

Conference attendees clearly welcomed the focus on applications, which led to animated discussions in the industry-government presentations. One workshop speaker took the tutorial by John Elder on common mistakes in data mining to heart; she did some real-time editing of her presentation to point out the mistakes in her application domain, such as a lack of caution in sampling the data and discounting pesky cases though they might reveal a larger problem in the data.

A new aspect of this year's conference was the increasingly important role of statistics in data mining. Keynote speaker Chris Bishop of Microsoft Research: Cambridge discussed recent advances in Bayesian inference techniques and several technical sessions focused on statistical techniques in data mining. This connection between statistics and data mining will be exploited further in the next conference in the series (scheduled for Newport Beach, April 21–23, 2005), which will be co-sponsored by the American Statistical Association and SIAM (http://www.siam.org/meetings/sdm05/). We encourage statisticians and data miners to submit papers and attend the conference, and help us narrow the gap between the two fields to bring together the best of both worlds.

The proceedings of SDM04, including the keynotes and the presentations at the industry/government session, are available on-line at http://www.siam.org/meetings/sdm04.

Department of Homeland Security Advanced Scientific Computing Workshops

Dates May 12, 2004 September 22–23, 2004

Locations Washington, DC Alexandria, Virginia

As part of the Department of Homeland Security (DHS) Advanced Scientific Computing program, ISCR co-hosted three requirements-gathering workshops during FY 2003. The first of these, the Advanced Scientific Computing Requirements Workshop was held October 8-9, 2003; a summary of this workshop can be found in the *ISCR 2003 Annual Report*. The two additional DHS workshops held this year were the Incident Management Simulation Workshop and the Data Sciences Workshop. The Krell Institute participated in the development of the content of these workshops and handled all workshop logistics and developed the workshop Web sites.

The Incident Management Simulation Workshop was held on May 12, 2004 at the Westin Grand Hotel in Washington, DC. This workshop brought together senior representatives of the emergency response and incident management communities with modeling and simulation technologists from DOE laboratories. The workshop provided an opportunity for incident responders to describe the nature and substance of the primary personnel roles in an incident response, identify current and anticipated roles of modeling and simulation in support of incident response, and begin a dialog between the incident response and simulation technology communities that will guide and inform planned modeling and simulation development for incident response.

The workshop was a joint effort of the Advanced Scientific Computing Program and the DHS Emergency Preparedness and Response Portfolio, both elements of the DHS Science and Technology Directorate. Based on the interactions at the workshop, a panel of computational science technologists prepared a summary report on incident management practice and the potential roles that computational simulation might play in supporting incident management. In particular, the panel prepared a summary of simulation capabilities that are relevant to incident management training and recommendations for the use of simulation in both incident management and in incident management training. In addition, the final report discusses areas where further research and development will be required to support future needs in this area.

The DHS Data Sciences Workshop was held September 22–23, 2004 at the Hilton Old Town in Alexandria, VA. The purpose of this workshop was to thoroughly review the data sciences mission needs of DHS and to identify specific mathematics and computer science research and development (R&D) topic areas required to address those needs. During the workshop, approximately 50 invited participants representing DHS, DOE and its national laboratories, academia and industry, identified specific R&D topic areas in the data sciences, their ties to the mission needs of DHS, and the potential impact of the proposed R&D. This effort will specify five years of relevant research topics in the data sciences area to support the Threat and Vulnerability Testing and Assessment Portfolio of DHS. These activities are important since they will immediately feed into the current DHS planning process for FY 2006.

Domain-Specific Languages for Numerical Optimization

Dates August 18–20, 2004

Location Argonne, Illinois

On August 18–20, 2004, ANL hosted a workshop on Domain-specific Languages for Numerical Optimization, co-sponsored by LLNL. There were 36 participants, including students, faculty, and staff from 12 universities, plus scientific staff from ANL, LLNL, and Sandia. The purpose of the meeting was to bring together experts in programming languages and compilers together with experts in numerical optimization and partial differential equations (PDEs) to stimulate discussion on the design and implementation of next-generation domain-specific languages for numerical optimization, with an emphasis on stochastic optimization and PDE-constrained optimization.

The participants discussed the design and implementation of current languages for numerical optimization, called modeling languages. These languages are mostly declarative, but take on an imperative flavor when a function or its derivatives must be evaluated. Native data types include scalars, sets, and ordered sets. Typically, the models and data are "compiled" into an internal representation or bytecode that is then interpreted. Several examples were given in which a problem that might have taken thousands of lines of code to express in Fortran or C required only 30-100 lines of code in AMPL or GAMS.

On the other hand, two examples were cited where converting a problem from a modeling language to C or Fortran (plus an FFT library in one case) resulted in a hundred- or thousand-fold speedup. Although the granularity of objects in modeling languages is typically much finer than that in other domain specific languages, it seems likely that static or dynamic compilation could provide the performance needed for some large problems without sacrificing the expressiveness of the modeling languages.

One of the obstacles to successfully extending modeling languages to support stochastic optimization is the wide variety of ways that randomness can enter an optimization problem and hence the many kinds of stochastic optimization problems. Even when the scope is restricted to a particular kind of stochastic optimization, multistage linear recourse problem, the specification of a problem is nontrivial. Part of the challenge arises from the fact that multistage problems can lead to enormous scenario trees. However, for problems in which stochasticity enters in a structured manner, extending a modeling language with time (stage) information and mechanisms to specify the probability distributions for random variables may suffice. The primary obstacle to extending modeling languages to support PDE-constrained optimization is that effective mechanisms for specifying partial differential equations themselves have not been developed. However, emerging systems, such as Sundance, FIAT, and PETSc 3, offer some hope that effective mechanisms for specifying and solving PDEs can be developed.

Several systems for analysis and transformation of general-purpose languages, domain-specific languages, and meta-languages were presented. The DMS Software Reengineering Toolkit supports automated source code analysis and modification. It utilizes Unicode lexers, GLR parsers for arbitrary context free grammars, analysis via multipass attribute grammars, and conditional source-to-source transformations. The extensible C (xtc) system uses a packrat parser to support arbitrary syntactic extensions to C, AST transformation rules to reduce and optimize, and typing rules to support safety constraints. Several researchers presented their work in the area of telescoping languages and related techniques. These methods exploit domain-specific analysis and optimization to improve the performance of general purpose languages extended with domain-specific libraries. Examples of systems supporting this paradigm are ROSE, Broadway, and libGen. In many cases, these systems are able to achieve performance superior to voluminous hand-developed Fortran or C implementations using concise implementations in C++ or MATLAB. It was also demonstrated that generic programming techniques can also provide high performance and high levels of expressiveness.

http://www-unix.mcs.anl.gov/workshops/DSLOpt/

Short-Pulse Laser Workshop

Dates August 25 – 27, 2004

Location Pleasanton, California

For three days at the end of August 2004, 55 plasma scientists met at the Four Points by Sheraton in Pleasanton to discuss some of the critical issues associated with the computational aspects of the interaction of short-pulse, high-intensity lasers with matter. The workshop was organized around the following areas of key interest to the Laboratory.

- Laser propagation / interaction through various density plasmas: micro scale.
- Anomalous electron transport effects: from micro to meso scale.
- Electron transport through plasmas: from meso to macro scale.
- Ion beam generation, transport, and focusing.
- "Atomic-scale" electron and proton stopping powers.
- Kα diagnostics.

Each area had a coordinator who drew up a list of questions, moderated discussions, and wrote a working group summary.

Many important problems in fast ignition are related to laser-plasma interactions, including laser propagation in the underdense corona plasma, laser hole-boring in the overdense plasma, laser absorption and energetic electron production at the critical surface, and electron transport in the mildly-dense plasma region. Participants worked out a set of benchmark computational simulation problems to compare their modeling capability in these areas.

Meso to macro scale electron transport discussions focused primarily on the correct method of initiating the electron beam. Several phenomenological techniques were discussed, such as injection at a plane in free space, promotion of ambient electrons, and the use of a ponderomotive force. It was generally agreed that the boundary conditions in the laser–plasma interaction (LPI) region were critical to the problem setup. A two-region approach, in which the LPI is simulated in the blow-off plasma and hybrid methods are used in the solid density material, might be a reasonable intermediate step. Since most codes do not have a laser-deposition package, it is necessary to choose and standardize the beam parameters, and a set was proposed.

The ion beam generation, transport, and focusing group discussed the following questions, from general to application specific.

- What are the proton generation mechanisms?
- What are their efficiencies?
- How sensitive to resolution are the answers?
- What codes can be used?
- How does electron flow affect proton generation?
- How can we control the generation and focusing of the protons?
- What is the optimum proton energy for radiography?
- What are the qualities that set ions using these mechanisms apart from "standard" ion beams?
- What governs ion flux?
- What is the optimal distance of the "proton lens" from the target?

The "atomic-scale" stopping powers session discussed first the stopping power of relativistic electron beams (REB) with energies of 1 to 10 MeV, stopping in pre-compressed deuterium–tritium (DT) targets, and the stopping power of non-relativistic (NR) protons with energies of 1 to 100 MeV. They set benchmark problems for multiple scattering of REB on target ions and of multiple scattering of NR protons in thin foils of high-Z materials disposed in front of laser proton sources (LPS). Finally, they considered REB and NR proton stopping in strongly magnetized fast ignition targets.

The purpose of the Kα diagnostic sessions was to discuss some of the progress being made in modeling Kα emission in short-pulse petawatt laser experiments and to discuss with experimentalists some of their latest results.

The workshop was made possible by the joint financial support of the Institute for Laser Science and Applications and the ISCR at LLNL.

Computational Methods in Transport

Dates September 11–16, 2004

Location Tahoe City, California

The Computational Methods in Transport Workshop was devoted to providing a forum where computational transport researchers could discuss their methods, successes and failures across disciplinary boundaries. Typically, the numerical methods used in a given field are communicated to other researchers in that field. Rarely, however, are those methods communicated between one application domain of radiation transport and another.

For example, nuclear engineers and astrophysicists rarely attend the same meetings or read the same literature. The goal of the Computational Methods in Transport Workshop was to address this discrepancy and open channels of communication and cooperation so that (1) existing methods used in one field could be applied to other fields and (2) greater scientific resources could be mobilized to help solve outstanding problems.

Beginning on the afternoon of September 11, 2004 and ending with lunch on September 16, 2004, the workshop was held at the Granlibakken Conference Center in Tahoe City, CA. The first day of the meeting consisted of a series of one-hour talks reviewing one of the major fields represented at the workshop. The areas covered included astrophysics, atmospheric physics, mathematics, plant canopies, nuclear engineering, oceanography and high-energy density physics. One afternoon was reserved for a poster session where 30 posters were presented in a very lively and well-attended event. The following days were filled with focused 45-minute talks by each of the representative fields that delved into more technical detail. Substantial time was reserved for individual networking and communications.

Speakers were chosen based on their international recognition and covered various topics:

- **Ed Larsen** (University of Michigan) — Numerical methods used in neutron transport.
- **Tony Mezzacappa** (ORNL) — Applications and numerical methods used in supernova core collapse.
- **David Levermore** (University of Maryland) — Moment and closure approximations used in approximating transport equations.
- **Marty Marinak** (LLNL) — Transport needs in high-energy density physics.
- **George Kattawar** (Texas A&M) – Polarization and radiative transfer in oceanography.

The conference structure and venue work extremely well. Participants from different fields, who would never have had the opportunity to speak with other participants, were engaged in stimulating and very fruitful discussions. The atmosphere was collegial with all participants willing to learn and teach. As the conference week progressed, atmospheric scientists were learning about methods used in nuclear engineering.

A radiation physicist related, "Using the Fokker–Planck equation for studying our scattering problems never occurred to me. We are going to look into this." A mathematician who gave a talk on medical imaging and radiation oncology forged a bond with a nuclear engineer and is taking a sabbatical leave to apply his knowledge to radiation oncology. Another outgrowth of the workshop is that Ryan Clement (LLNL) and organizer Frank Graziani are setting up an e-print server that will serve as a repository for computational transport papers.

Overall, feedback from the workshop has been very positive and most expressed hope that it would be done again. The plan is to do a smaller, multi-disciplinary special-topics meeting next year followed by a larger meeting in 2006 similar to the 2004 workshop. Like the 2004 workshop, the future ones will be organized in conjunction with the Institute for Pure and Applied Mathematics at UCLA.

http://www.ipam.ucla.edu/programs/tr2004/

AMG / FOSLS Summit

Dates September 27 – October 3, 2004

Location Lake City, Colorado

In 1997, an informal meeting was held in Frisco, CO, between researchers in CU Boulder's Applied Mathematics Department and LLNL's CASC Division to discuss their collaboration on Algebraic Multigrid (AMG) methods. They met again in Boulder in 1998. In 2000, the meetings became annual, held in Lake City, CO. In 2002, they were expanded to include discussions on the First Order Systems Least Squares (FOSLS) methodology.

The summit is structured as a "working meeting" with an emphasis on exposing open research issues and generating ideas for solving them. Formal talks are strongly discouraged in favor of whiteboard discussions and individual interactions, a format that distinguishes it from typical meetings and conferences. Participation is by invitation only, consisting primarily of researchers from CU Boulder and CASC, but also including a small number of leading experts from other institutions around the world.

The Summit was held September 27–October 3, 2004. The first half of the meeting focused on AMG, and the second half focused on FOSLS, with an overlap day in between. There were 31 attendees this year. The main CU/CASC group consisted of 11 from CU Boulder, 9 from CASC, plus 4 recent CU graduates. The other attendees were: Irad Yavneh (Technion, Israel), Ludmil Zikatanov (Penn State), Ira Livshits (Ball State), Achi Brandt Weizmann Institute, Israel), Bobby Philip (LANL), Marzio Sala and Michael Gee (SNL).

The first half hour of the meeting was spent setting the agenda. The topics suggested this year by the attendees were as follows (as written on the whiteboard).

- Weighted Functionals
- Smooth Aggregation & e-Free AMGe
- Nonlinearity
- Almost Zero Modes
- Measures
- Coarse Variable Types
- Sparsity of P/Dilution
- Relaxation
- Sharp Theory
- Trace Minimization
- Iradism
- Upscaling
- Wavelet AMG

During the remainder of the meeting, each topic item (and its associated issues) was discussed in detail and solution approaches were proposed and debated. Two sample outcomes of the meeting were an improved theoretical foundation for a new trace minimization approach for defining interpolation in AMG, and an idea for relating a new compatible relaxation method (one that defines the coarse variables as averages) to a recent AMG theory and framework.

Summer Research Program

The ISCR put on its largest summer visitor program ever in FY 2004 with approximately 100 students in residence, as well as 11 faculty. Students are assigned individual LLNL mentors and given specific projects, ranging from programming tasks to original research, to which they will contribute based on their background and skills. Reports for most of the projects in the following table have been omitted from the printed version of the ISCR Annual Report in the interest of space, but can be found on the accompanying PDF. You can also obtain a CD-ROM containing these reports by calling (925) 423-3691.

To view the full abstract of any of the seminars listed on the following pages, simply click on the name of the seminar.

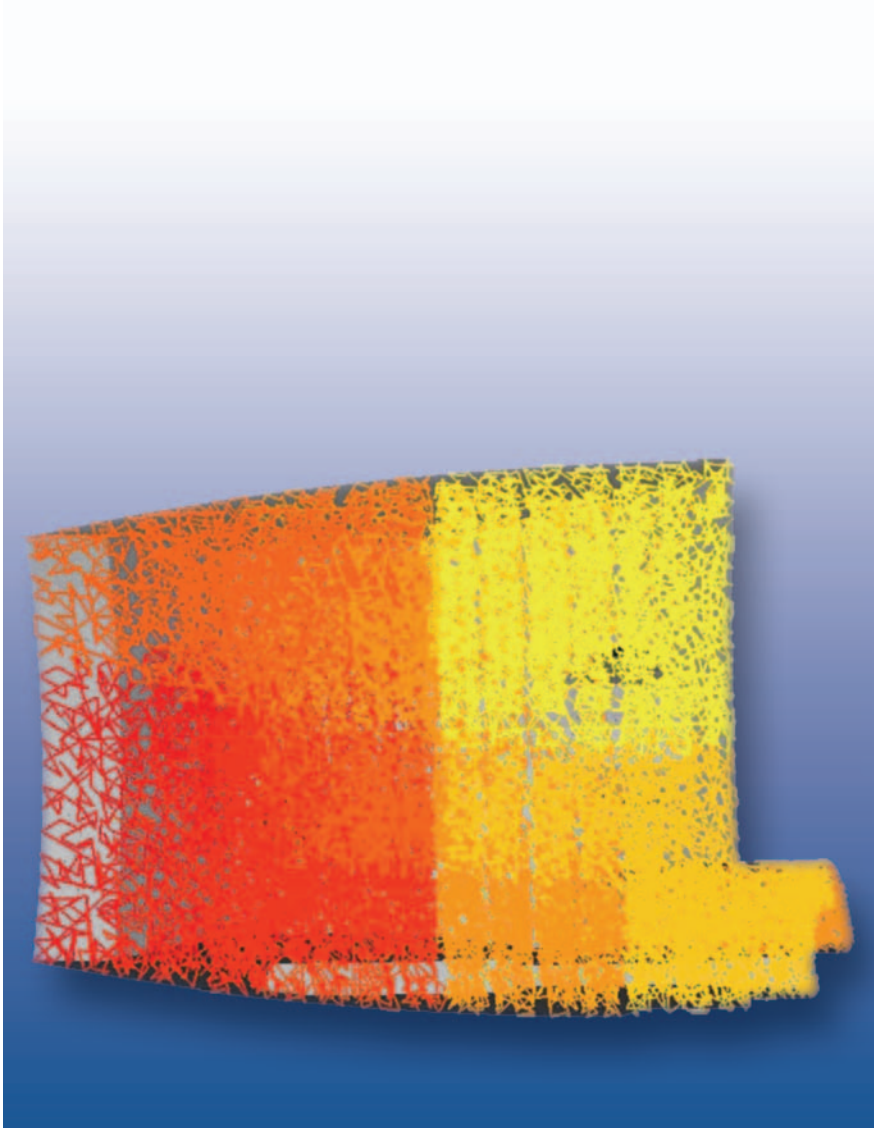

TOC

Summer Research Program

Name, Affiliation	Mentor, Host Organization
Project Title	

Sabbatical Visitors

Carl Ollivier-Gooch, University of British Columbia	Lori Diachin, CASC
Unstructured Mesh Technology Development	
Anne Greenbaum, University of Washington	Peter Brown, CASC
Boundary Integral Equation Methods and Software for Poisson's Equation in Three Dimensions	
Bjorn Sjogreen, Royal Institute of Technology, Stockholm, Sweden	Anders Petersson, CASC
Embedded Boundary Methods for Partial Differential Equations	

Summer Faculty

Constantin Bacuta, University of Delaware	Rob Falgout, CASC
A New Approach for Solving the Stokes Problem Based on the Distributed Relaxation Method	
James Brannick, University of Colorado at Boulder **Ludmil Zikatanov,** Pennsylvania State University	Rob Falgout, CASC Panayot Vassilevski, CASC
Compatible Relaxation and Optimal AMG Interpolation	
Zhiqiang Cai, Purdue University	Panayot Vassilevski, CASC
Mixed and Multigrid Methods for Stokes and Navier–Stokes Equations	
Tim Chartier, Davidson College	Van Emden Henson, CASC
Adaptive Multigrid via Subcycling on Complementary Grids	
Irene Livshits, University of Central Arkansas	Rob Falgout, CASC
AMG Algorithm for Finding an Eigenbasis for the Schroedinger Operator	
Anne Ngu, Southwest Texas University	Terence Critchlow, CASC
Large-Scale Integration of Web Sources	

Summer Students

David Alber, University of Illinois	Jim Jones, CASC
Adapting Algebraic Multigrid for the Solution of the Curl–Curl Formulation of Maxwell's Equations	
Kristopher Andersen, Unversity of California, Davis	John Pask, PAT
Locally Optimal Methods to Solve Eigenvalue Problems in Electronic Structure Calculations	
John Anderson, University of California, Davis	Benjy Grover, DCOM
QTester	
Dustin Anderson, California Polytechnic State University, San Luis Obispo	Don MacQueen, EPD
Real-Time Radiation Area Monitoring: Emergency Response and Regulation	
Benjamin Apodaca, Northern Arizona University	Carolyn Wimple, NIFE
PIMS Regression Test Development and Product Enhancement	

Name, Affiliation	Mentor, Host Organization
Project Title	
Abraham Bagherjeiran, University of Houston	Erick Cántu–Paz, CASC
Design and Implementation of an Anomaly Detector	
Paul Baginski, University of California, Berkeley	Van Emden Henson, CASC
Scalable Graph Algorithms	
Jeffrey W. Banks, Rensselaer Polytechnic Institute	William Henshaw, CASC
High-Speed Multicomponent Flows	
Lerone Banks, University of California, Davis	Jerry Rayome, CSP
Firewall Egress Filtering	
Janine Bennett, University of California, Davis	Valerio Pascucci, CASC
Volumetric Mesh Parameterization towards Slow-Growing Subdivision	
James Brannick, University of Colorado at Boulder	Rob Falgout, CASC
Adaptive Algebraic Multigrid Preconditioners in Quantum Chromodynamics	
Steve Callahan, University of Utah	Valerio Pascucci, CASC
A Memory-Insensitive Format for Out-of-Core Access to Unstructured Volumetric Meshes	
Jedidiah Chow, Granada High School	Jean Shuler, CSG
Discovery Center Display	
Kevin Chu, Massachusetts Institute of Technology	David Trebotich, CASC
Incorporating Electrokinetic Effects into the EB Navier–Stokes Embedded Boundary Incompressible Fluid Solver	
Kree Cole-McLaughlin, University of California, Los Angeles	Valerio Pascucci, CASC
Applying Morse Theory to Computational Datasets	**Terence Critchlow, CASC**
Dylan Copeland, Texas A & M University	Barry Lee, CASC
Geometric Multigrid for Variable Coefficient Maxwell's Equations	
Steven Davis, California State University, Hayward	Charles Doutriaux, EEBI
Web Designer, Program for Climate Model Diagnosis and Intercomparison	
Stanko Dimitrov, University of Arizona	Edmond Chow, CASC
Edge Betweenness Properties in Complex Networks	
Veselin Dobrev, Texas A & M University	Panayot Vassilevski, CASC
Element Agglomeration AMGe Solvers for Unstructured Finite-Element Problems	
Nina Dokeva, University of Southern California	Panayot Vassilevski, CASC
Parallel Implementation of a FETI–DP Method for Elliptic Problems with Mortar Finite Element Discretization	
Emily Eder, University of California, Los Angeles	Robert Fernandes, IOAC
Graph Viewer Improvements	

Summer Research Program

Name, Affiliation	Mentor, Host Organization
Project Title	
Christopher Egner, Rochester Institute of Technology	Paul Amala, AX Division
Platform Independence for Cale: wxWidgets and GNU Autotools	
Chris Elofson, University of Arizona	Vijay Sonnad, DCOM
Wavelet-Based Opacity Data Compression	
Dayo Esho, University of California, Berkeley	Steven Lee, CASC
ODE Visualization and Archive Tool in Python	
Pak-Wing Fok, Massachusetts Institute of Technology	Petri Fast, CASC
2D Numerical Modelling of Soap Film in Overture	
Camille Fournier, University of Wisconsin-Madison	Michael Kumbera, DCOM
Implementing Cray Pointers in the GNU gFortran Compiler	
Jeffrey L. Freschl, University of Wisconsin-Madison	John Johnson, DCOM
Performance Analysis of Monitoring and Information Systems Using NetLogger and an Investigation into P2P Overlay Topologies	
Evan Geller, Summit Preparatory High School	Albert Chu, HPSD
Bringing Order to CHAOS	
Tobias Gradl, Technische Universität München, Germany	Edmond Chow, CASC
Assessing Performance of Hybrid MPI/OpenMP Programs on SMP Clusters	
Rachel Greenstadt, Harvard University	David Youd, NAIC
Censorship Resistance	
Attila Gyulassy, University of California, Davis	Valerio Pascucci, CASC
Hierarchical Morse–Smale Complexes	
Daniel Han, University of Southern California	Terry Jones, ICC Services and Development Division
MPI Profiling	
Kevin Hoffman, Brigham Young University	Robert Cooper, DCOM
Parallel Analysis of Asymmetries in Symmetric Simulations in ALE3D	
W. Taylor Holliday, University of California, Davis	Valerio Pascucci, CASC
Combinatorial Feature Extraction for a Streaming Framework	
Gary Hon, University of California, San Diego	David Clague, EE-EETD
Virtual PCR (vPCR)	
Jason Howell, Clemson University	Carol Woodward, CASC
A Two-Grid Method for Radiation Diffusion	
David Hoyt, Brigham Young University	James Schek, NAIC
DHS Countermeasures Test Bed–Command Center	

Name, Affiliation	Mentor, Host Organization
Project Title	
Kevin Huck, University of Oregon	Brian Miller, CASC
PerfTrack	
Lukas Jager, Universität Bonn, Germany	Radu Serban, CASC
A Multilevel Preconditioner for a PDE-Constraint Optimization Problem	
Abram Jujunashvili, University of Colorado at Denver	Charles Tong, CASC
Integration of Multivectors to Hypre	
Alin Jula, Texas A&M University	Dan Quinlan, CASC
Automatic Detection of Hot Spots in C/C++ Programs	
Kirk Kelsey, University of Rochester	Tom Epperly, CASC
Automatic SIDL Generation with ROSE	
Jason Kimball, University of California, Irvine	Mark Duchaineau, CASC
Large-Scale Atom Data Visualization	
Nicholas Kridler, University of Colorado at Boulder	Barry Lee, CASC
Algebraic Multigrid for Maxwell's Equations with Variable Coefficients	
Brian Lum, University of California, San Diego	Amber Marsh, EEBI
IMAGE Pipeline	
Eric Machorro, University of Washington	Britton Chang, CASC
Numerical Solution Methods and Error Analysis for the Neutron Transport Equation in Absorbing, Monoenergetic, Non-Scattering Media	
Anna Majkowska, University of California, Riverside	Valerio Pascucci, CASC
Out-of-Core Visualization of Climate Modeling Data	
Robert Dean Malmgren, Northwestern University	Scott Brandon, AX Division
Code-Independent Analysis Tools for Physics Simulation Codes	
Peter Manning, University of Arizona	Glenn Goderre, NIFE
Development of a Graphical User Interface for the Virtual Beam Line (VBL) Simulation Software	
Ajith Mascarenhas, University of North Carolina at Chapel Hill	Valerio Pascucci, CASC
Time-Varying Reeb Graphs for Space-Time Data	
Tammara Massey, University of California, Los Angeles	Terry Brugger, NAIC
Ultra-Wideband Token Ring Simulation and Security in Sensor Networks at LLNL	
Kathryn Mohror, Portland State University	John May, CASC
PerfTrack	
Christopher Muelder, University of California, Davis	Marv Christensen, NAIC
Summarizing Network Traffic with Information Visualization	

Summer Research Program

Name, Affiliation	Mentor, Host Organization
Project Title	
Spencer Nielsen, Brigham Young University	Darrel Whitney, CADSE
Software Failure Risk Analysis	
Daniel Oeltz, Universität Bonn, Germany	Panayot Vassilevski, CASC
A Space-Time Sparse Grid Approximation Space	
Christopher Olson, University of California, Santa Cruz	Robert Cooper, DCOM
Reviewing Vista	
Timothy Paik, University of California, Berkeley	Shawn Larsen, ISCR Dave Harris, ISCR
Subspace Detectors	
Sung W. Park, University of California, Davis	Peter Lindstrom, CASC
Streaming Pointsets	
Bryan Parno, Harvard University	Tony Bartoletti, NAIC
Internet Ballistics: Retrieving Forensic Data From Network Scans	
Dan Phung, Columbia University	Moe Jette, CASC
SLURM for BlueGene/L	
Sriram Polepeddi, Carnegie Mellon University	Noel Tijerino, IOAC
Software Vulnerability Taxonomy Consolidation	
Raymond Pon, University of California, Los Angeles	Terence Critchlow, CASC
Performance-Oriented, Privacy-Preserving Data Integration	
Serban Porumbescu, University of California, Davis	Mark Duchaineau, CASC
Out-of-Core Parameterization and Remeshing of Surfaces	
Ashley President, Carnegie Mellon University	Michael Kumbera, A Division
Intrinsic Function Testing and an OpenMP Runtime Library Summer	
Davinder Rama, California State University, Sacramento	Adam Zemla, EEBI
Amino-Acid Sequence into Tertiary Structure Web Site	
George Roberts, Georgia Institute of Technology	Chandrika Kamath, CASC
Variance-Based Feature Tracking	
Rolf Ryham, Penn State University	Rob Falgout, CASC
Multigrid Prolongation Based on Sharp Convergence Theory	
Andreas Saebjornsen, Universitete i Oslo, Norway	Dan Quinlan, CASC
The AST Query Mechanism and the C/C++ Graphing Mechanism	
Elmer Salazar, California State University, Stanislaus	Michael Kumbera, DCOM
Implementation of OpenMP Support in the gFortran Compiler	
Jennifer Sirp, California State University, Sacramento	Terry Brugger, NAIC
Extending ReiserFS for Automatic File Queuing	

Name, Affiliation	Mentor, Host Organization
Project Title	
Yoshifumi Suzuki, University of Michigan	Jeff Hittinger, CASC
Efficient Schemes for Hyperbolic Systems with Stiff Relaxation Source Terms	
Valerie Szudziejka, University of California, Davis	Valerio Pascucci, CASC
Streaming Computation of Structural Graphs	
Ryan Szypowski, University of California, San Diego	Ulrike Yang, CASC
Parallel AMG for Systems of PDEs	
Brian Taylor, University of Illinois at Urbana-Champaign	Bill Henshaw, CASC
Computation of Cellular Detonation	
Robert Taylor, Northern Arizona University	Jody Malik, CSP
Vulnerability Tracking Database 2.0	
Nils Thuerey, Universität Erlangen-Nürnberg, Germany	Dan Quinlan, CASC
Global Analysis of the ROSE Infrastructure	Jody Malik, CSP
Peter Tipton, University of Southern California	Brandon Scott, AX Division
Code Validation Made Easy	
Chunbo Wang, Purdue University	Charles Tong, CASC
A MATLAB Implementation of Mixed Finite-Element Method for Incompressible Newtonian Flows: Pseudostress–Velocity Formulation	
Rebecca Wasyk, Worcester Polytechnic Institute	Carol Woodward, CASC
Newton–Krylov Methods for Expensive Nonlinear Function Evaluations	
Sage Weil, University of California, Santa Cruz	Tyce McLarty, CADSE
Metadata Management for Petabyte-Scale File Systems	
Dan Wendlandt, Stanford University	Martin Casado, NAIC
Information Leakage Due to Geographic Properties of Internet Routing	
Brian White, Cornell University	Dan Quinlan, CASC
Ameliorating the Performance Degradation of User-defined Abstractions and Indirect Memory Accesses	
Ryan M. White, University of California, Berkeley	Shawn Newsam, CASC
Matching Shapes Using Local Descriptors	
Jeremiah Willcock, Indiana University	Dan Quinlan, CASC
Additions to the ROSE Compiler Infrastructure	
Suzanne Wingenter, San Diego State University	Petri Fast, CASC
Shallow-Water Equations on Curvilinear Grids	
Christopher Wojtan, University of Illinois, Urbana-Champaign (2000–2004) Georgia Institute of Technology (2004)	Jeremy Meredith, DNT
A Hybrid Sort-First/Sort-Last Approach for Rendering Translucent Geometry in the VisIt Visualization Tool	

Summer Research Program

Name, Affiliation	Mentor, Host Organization
Project Title	
Sung-Eui Yoon, University of North Carolina at Chapel Hill	Peter Lindstrom, CASC
Cache Coherent Mesh Layout	
Beth Yost, Virginia Polytechnic Institute and State University	Terence Critchlow, CASC
A Visual Interface for the Promoter Identification Workflow	

ISCR Seminar Series

The ISCR hosted 70 seminars from visitors in FY 2004 covering a wide spectrum of research areas, and recruited an additional 35 speakers from LLNL ranks to speak to visiting students. The ASC Institute for Terascale Simulation Lecture Series was established in 2000 to enrich the intellectual atmosphere of LLNL's large simulation community through the visits of leaders representing the diverse areas of computation. In FY 2004, we hosted William Wulf, John Grosh, Bjarne Stroustrup, Mary Wheeler, and David Bailey. The ISCR Summer Student Lecture Series was also established in 2000 and forked into three different series in Summer 2004—Internships in Computational Modeling at the Terascale (ICMT), Internships in Computer Science at the Terascale (ICST), and College Cyber Defenders Computer Security. Seminar reports listed in the following table have been omitted from the printed version of the ISCR Annual Report in the interest of space, but can be found on the accompanying PDF. You can also obtain a CD-ROM containing these reports by calling (925) 423-3691.

To view the full abstract of any of the seminars listed on the following pages, simply click on the name of the seminar.

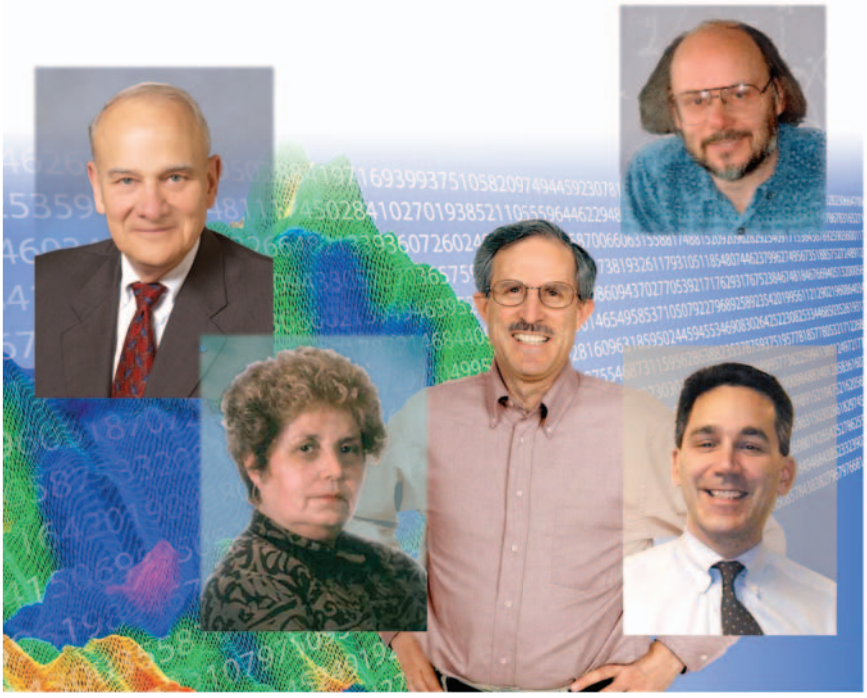

TOC

ISCR Seminar Series

Date	Speaker, Affiliation Title of Seminar
ITS Lecture Series	
12/4/03	**William Wulf,** National Academy of Engineering Challenges for Computing and Information Technology in the 21st Century
2/17/04	**John Grosh,** Department of Defense A DoD Perspective on High-End Computing
3/1/04	**Bjarne Stroustrup,** Texas A&M University Programming, Language and Libraries
6/24/04	**Mary Wheeler,** University of Texas at Austin Mathematical and Computational Modeling of Multiphysics Couplings
8/4/04	**David Bailey,** Lawrence Berkeley National Laboratory Twelve Ways to Fool the Masses: Back to the Future
Visitor Seminars	
10/1/03	**Kenneth Chiu,** Indiana University Bloomington A C++ Reflection Library
10/3/03	**Branden Fitelson,** University of California, Berkeley A User-Friendly Probability Machine with Applications
10/9/03	**Olav Beckmann,** Imperial College London,United Kingdom Dynamic Code Generation in C++ as a Foundation for Domain-Specific Optimization
10/23/03	**Charles Elkan,** University of California, San Diego Progress in Clustering
10/23/03	**James Glimm,** Stony Brook University Statistical Models for Simulation Errors and Their Role in Prediction and Uncertainty Quantification
10/24/03	**Xiaolin Li,** Stony Brook University Simplification, Adaptivity and Conservation in Front Tracking Method
10/27/03	**Achi Brandt,** Weizmann Institute of Science, Israel Image Processing: From Segmentation to Recognition
10/30/03	**Douglas Swesty,** Stony Brook University Efficient Solution of the Discretized S_N Time-Dependent Boltzmann Transport Equation on Parallel Platforms
10/31/03	**Frank Mueller,** North Carolina State University Detailed Cache Coherence Characterization for OpenMP Benchmarks
11/7/03	**Daniel Segre,** Harvard Medical School Computational Models of Cellular Metabolic Fluxes

Date	Speaker, Affiliation Title of Seminar
11/13/03	Richard Scalettar, University of California, Davis Linear Algebra and Quantum Simulations
11/14/03	Marsha Berger, New York University Cartesian Grids with Embedded Geometry
11/24/03	Johannes Gehrke, Cornell University Techniques for Processing Large Data Streams
12/3/03	John Feo, Cray Inc. Cray Cascade Project
12/3/03	Ravi Samtaney, Princeton Plasma Physics Laboratory The Magneto–Hydrodynamic Richtmyer–Meshkov Instability
12/9/03	Margot Gerritsen, Stanford University Why Are Streamline Methods Attractive for Simulation of Gas-Injection Processes?
12/9/03	Jonas Nilsson, Stanford University A Hybrid High-Order Method for the Incompressible Navier–Stokes Equations
12/12/03	Barney Maccabe, University of New Mexico Early Experience in Splintering Communication Protocols
12/12/03	Jennifer Widom, Stanford University The Stanford Data Stream Management System
12/17/03	Daniel Crichton, Jet Propulsion Laboratory A Data Grid Framework for Managing Planetary Science Data
12/17/03	William Dally and Patrick Hanrahan, Stanford University Merrimac: Supercomputing with Streams
12/18/03	Terran Lane, University of New Mexico From Security to Cells: Ongoing Machine Learning Research at the University of New Mexico
1/9/04	Christopher Jermaine, University of Florida Approximate Query Processing with Sampling and Pre-Aggregation
1/12/04	Thomas Seidman, University of Maryland, Baltimore County Hybrid Systems: Discontinuous Dynamics in a Continuous World
1/16/04	Serge Belongie, University of California, San Diego Three Brown Mice: See How They Run — Monitoring Rodent Behavior in the Smart Vivarium
1/23/04	Cheryl McCosh, Rice University Type-Based Specialization in a Telescoping Compiler for MATLAB
1/26/04	Elizabeth Post, Lincoln University On the Farm: Parallel Small Talk for Simulating Dairy Operations

ISCR Seminar Series

Date	Speaker, Affiliation Title of Seminar
1/27/04	**Steven Parker,** University of Utah What's New with SCIRun2?
1/29/04	**Miguel Argaez,** University of Texas at El Paso An Optimization Technique for Large-Scale Nonlinear Programming
1/29/04	**Leticia Velazquez,** University of Texas at El Paso A Global Optimization Technique for Solving Zero or Very Small Residual Nonlinear Least-Squares Problems
2/10/04	**James Hobart,** Classic Systems Solutions Designing for Usability
2/10/04	**Alan Laub,** University of California, Davis Statistical Condition Estimation
2/11/04	**Demet Aksoy,** University of California, Davis PLASMA (PLAnetary Scale Monitoring Architecture)
2/18/04	**Donald Schwendeman,** Rensselaer Polytechnic Institute Numerical Method for High-Speed Reactive Flow on Overlapping Grids
2/20/04	**Matteo Pellegrini,** University of California, Los Angeles PROLINKS: A Database of Co-Evolving Proteins
2/25/04	**Laxmikant Kale,** University of Illinois, Urbana-Champaign Adaptive Resource Management via Processor Virtualization: Charm++ and AMP
3/3/04	**Lada Adamic,** Hewlett-Packard Laboratories How to Search a Social Network
3/4/04	**Christoph Pflaum,** Universität Erlangen-Nürnberg, Germany 3D Computation of Laser Cavity Eigenmodes by Finite Elements
3/5/04	**Joel Saltz,** Ohio State University Middleware Support for Data Ensemble Analysis
3/10/04	**Stephen Neuendorffer,** University of California, Berkeley Actor-Oriented Metaprogramming
3/15/04	**Luiz DeRose,** IBM Research DPOMP: An Infrastructure for Performance Monitoring of OpenMP Applications
3/18/04	**Boleslaw Szymanski,** Rensselaer Polytechnic Institute Analyzing Evolution of Virulence through Spatially-Explicit Epidemic Models
3/26/04	**Gunther Weber,** University of California, Davis Topology-Based Exploration of Scalar Fields

Date	Speaker, Affiliation Title of Seminar
4/6/04	**Kirk Hays and Max Alt,** Intel Corporation Performance Analysis and Tuning of MDCASK and PF3D Codes on Itanium Processors
5/6/04	**David Liu,** University of California, Berkeley GridDB: A Data-Centric Overlay for Scientific Grids
5/14/04	**Scott Baden & Jacob Sorenson,** University of California, San Diego Data-Driven Execution of Communication Tolerant Algorithms
5/17/04	**Homer Walker,** Worcester Polytechnic Institute Globalization Techniques for Newton–Krylov Methods
5/24/04	**Jeffrey Heys,** University of Colorado at Boulder Numerical Issues When Modeling Fluid-Elastic Interaction in 3D with First-Order System Least Squares
5/24/04	**Chad Westphal,** University of Colorado at Boulder First-Order System Least-Squares for Problems with Boundary Singularities
6/1/04	**Martin Bazant,** Massachusetts Institute of Technology Induced-Charge Electro-osmosis
6/3/04	**Marc Barthelemy,** Commissariat á l'Energie Atomique, France Structure and Modeling of Weighted Complex Networks
6/8/04	**David Jensen,** University of Massachusetts Amherst Knowledge Discovery in Networks
6/10/04	**Steven Knight,** SCons Project SCons: A Next-Generation Build Tool
6/16/04	**Wenke Lee,** Georgia Institute of Technology Worm Detection and Response: Local Strategies and Analytical Models
7/21/04	**Frank Mueller,** North Carolina State University Detailed Cache Coherence Characterization for OpenMP Benchmarks
7/23/04	**Sameer Agarwal,** University of California, San Diego On Refractive Optical Flow
7/23/04	**Gabriele Jost,** NASA Ames Research Center What Multilevel Parallel Programs Do When You Are Not Watching
7/29/04	**David Forsyth,** University of California, Berkeley Words and Pictures
7/29/04	**Ling Liu,** Georgia Institute of Technology ReFlex: Flexible and Reliable Systems Technologies for Responding to Massively Disruptive Events

ISCR Seminar Series

Date	Speaker, Affiliation Title of Seminar
7/29/04	Nikolaos Nikiforakis, Cambridge University, United Kingdom Emergence of Detonation in the Flowfield Induced by Richtmyer-Meshkov Instability
8/13/04	Chris Wiggins, Columbia University Data-Driven Approaches for Biological Networks: Inference, Organization and Analysis
8/27/04	Alex Schweitzer, Universität Bonn, Germany Efficient Implementation and Parallelization of Meshfree and Particle Methods: The Parallel Multilevel Partition of Unity Method
9/8/04	E. Ann Stanley, Los Alamos National Laboratory Using Mathematical Models to Understand the AIDS Epidemic and Guide Policy
9/9/04	Jia Li, University of Alabama, Huntsville Mathematical Modeling of Malaria, Early Warning System, and Transgenic Mosquitoes
9/30/04	Reagan Moore, San Diego Supercomputer Center Digital Libraries and Data-Intensive Computing

LLNL Summer Seminars, ICST Computer Science

6/15/04	Pat Miller, CASC FlashMob Instant Supercomputing
6/22/04	David Jefferson, CASC The Time Warp Method of Parallel Discrete Event Simulation.
6/29/04	Kim Yates, CASC BlueGene/L: The Next Generation of Scalable Supercomputer
7/13/04	Chandrika Kamath, CASC Scientific Data Mining: The Sapphire Project
7/20/04	Tom Epperly, CASC Babel Language Interoperability Tool
7/27/04	Gregory M. Pope, CADSE Why Software Quality Assurance Practices Become Evil!
8/3/04	Kim Minuzzo, NIF NIF Control System
8/10/04	Valerio Pascucci, CASC Multiresolution Computation and Presentation of Topological Structures
8/17/04	Erick Cantú-Paz, CASC Solving Problems with Evolutionary Algorithms

Date	Speaker, Affiliation Title of Seminar

LLNL Summer Seminars, ICMT Computational Science

6/9/04	David Keyes, Columbia University and ISCR A Science-Based Case for Large-Scale Simulation
6/16/04	Gary Kumfert, CASC Power Presentations
6/30/04	Peter G. Eltgroth, CASC Overview of Research Activities in the Center for Applied Scientific Computing
7/6/04	Paul Dubois, CASC Effective Architectures for Scientific Programs
7/7/04	Steven F. Ashby, CAR Computational Science at Lawrence Livermore National Laboratory
7/14/04	David Trebotich, CASC Computational Microfluidics
7/21/04	Frank Graziani, B Division Radiation Transport in 3D Random Media: Direct Numerical Simulation
7/28/04	Don Dossa, CASC Architectural Overview of BlueGene/L Supercomputer
8/11/04	Rose McCallen and Kambi Salari, CASC DOE's Effort to Reduce Truck Aerodynamic Drag — Joint Experiments and Computations Lead to Smart Design
8/18/04	Lori Freitag-Diachin, CASC Developing Interoperable Meshing and Discretization Components
8/25/05	Rob Falgout, CASC Scalable Linear Solvers: Multigrid Methods

LLNL Summer Seminars, CCD Computer Security

6/10/04	Terry Brugger, NAIC Data Mining for Network Intrusion Detection
7/1/04	Tony Bartoletti, NAIC Identifying Internet Adversaries Despite Falsified Source Addressing
7/8/04	Bill Orvis, EE-EETD The Tale of Two Rootkits — SucKit and Hacker Defenders

Acronyms and Abbreviations

1D 2D 3D AAWP AC AIDS ALC ASC ALE3D AMG aAMG AMGe AMPI AMPL AMR ANL AP API ARPACK ASC AST AS2TS BELCh BGK BGL BLAS BLAST BMP BOOST BSP CADSE CAR CASC dbCAT CCA CCD CES CFD CFL CHAOS CHOMBO CIAC CJ COCs CORBA CR CRON C-SAFE CSGF CSS CSV CU CVDB DAKOTA DAG DARPA DC DCOM DDT DG DHS DMS DMTs DoD DNA DNS DNS DOE DPCL DPIV DPOMP DSCEAs EC EE EETD EEBI EFGM EOC EPIC ERD ERI ESS EGG FAS FASTA FBA FEM FETI-DP FFLO FFT FIAT FORTRAN FOSPACK FOSLS FY GAMS GCC GEM GHHE GLR GMRES GNU GridDB GUI HCI HIV HkDef HLL HP HPC HTML HTTP HVAC HYDRA HYPRE IBM ICAT ICCD ICCS ICEO ICMT ICST IDS ISO IDX ILP IMAGE IMPACT IOR IP IPAM I/O ISCR ISMG IT ITER ITS JDO JPEG JSP KINSOL KOJAK KULL LANL LDA LDRD LES LGA LINPACK LLNL LOBPCG LOCKSS LOD LPI LPS MATLAB MCR MDCASK MEMS METIS MHD MIRANDA MLIC MOMA MPI MRF fMRI NAI NAIC NAMD NASA NCAR NFS NIF NIFE NNSA NR ODE OODT OpenMP ORNL OS OSVDB PAPI PAT PaToH PC PCG PCMDI PCR vPCR PDB PDEs PDS PERC PerfTrack PF3D PHP PHP/MySQL PI PIMS PIW PLASMA PMaC PMPI POMP POP POSIX PROLINKS PUM PVODE QBF QC QCD RANS RANSAC RBF REB

TOC

Acronyms and Abbreviations

A

1D – one-dimensional

2D – two-dimensional

3D – three-dimensional

AAWP – Analytical Active Worm Propagation

AC – alternating current

AIDS – acquired immune deficiency syndrome

ALC – ASC (see) Linux Cluster, provides computing cycles for ASC Alliance users and unclassified ASC code development. ALC and MCR (see) are sibling Linux clusters with ~1000 Intel P4 nodes with dual 2.4GHz processors.

ALE3D – Arbitrary Lagrangian-Eulerian 3D code

AMG – Algebraic Multigrid method developed to solve matrix equations resulting from the discretization of an elliptical PDE on an unstructured mesh

aAMG – adaptive AMG

AMGe – Algebraic Multigrid based on element interpolation

AMPI – Adaptive Message Passing Interface

AMPL – A Mathematical Programming Language

AMR – adaptive mesh refinement

ANL – Argonne National Laboratory

AP – asymptotic preservation

API – Application Program Interface

ARPACK – a collection of Fortran77 subroutines designed to solve large-scale eigenvalue problems

ASC – Advanced Simulation and Computing

AST – abstract syntax tree

AS2TS – Amino-acid sequence into tertiary structure

B

BELCh – Boundary Exponential Log Characteristic

BGK – Bhatnagar-Gross-Krook

BGL – BlueGene/L, an IBM computer, currently the world's fastest

BLAS – Basic Linear Algebra Subprograms

BLAST – Basic Local Alignment Search Tool

BMP – file extension for bitmap graphic files

BOOST – C++ libraries aimed at providing quality software components to developers using Standard Template Library styles.

BSP – bulk synchronous parallelism

C

CADSE – Center for Applications Development & Software Engineering

CAR – Computing and Applied Research

CASC – Center for Applied Scientific Computing

dbCAT – Catalog of Databases

CCA – Common Component Architecture

CCD – College Cyber Defenders

CES – Cost Effective Sampling

CFD – Computational Fluid Dynamics

CFL – Courant-Friedrichs-Lewy

CHAOS – Clustered High-Availability Operating System

CHOMBO – a set of software tools for implementing finite difference methods for the solution of PDEs (see) on block-structured adaptively refined rectangular grids

CIAC – Computer Incident Advisory Capability

CJ – Chapman-Jouguet

COCs – chain of custodies

CORBA – Common Object Request Broker Architecture

CR – compatible relaxation

CRON – not an acronym

C-SAFE – Center for Simulation of Accidental Fires and Explosions

CSGF – Computational Science Graduate Fellowship

CSS – Cascading Style Sheets

CSV – Comma-Separated Values

CU – University of Colorado

CVDB – consolidated vulnerability database

D

DAKOTA – Design Analysis Kit for Optimization and Terascale Applications

DAG – Directed Acyclic Graph

DARPA – Defense Advanced Research Projects Agency

DC – direct current

DCOM – DNT and PAT Computing Applications Division

DDT – deflagration to detonation transition

DG – Discontinuous Galerkin

DHS – Department of Homeland Security

DMS – Design Maintenance Systems

DMTs – data management tools

DoD – Department of Defense

DNA – deoxyribonucleic acid

DNS – direct numerical simulation

DNS – domain name system

DOE – Department of Energy

DPCL – Dynamic Probe Class Library

DPIV – Digital Particle Image Velocimetry

DPOMP – Dynamic Performance Monitor for OpenMP

DSC – Destination Source Correlation

E

EAs – evolutionary algorithms

EC – exponential characteristic

EE – Electronics Engineering

EETD – EE Technologies Division

EEBI – Energy & Environment, Biology & Biotechnology, and Institutional Computing

EFGM – element-free Galerkin methods

EOC – Emergency Operations Center

EPIC – Explicitly Parallel Instruction Computing

ERD – Environmental Restoration Division

ERI – Exploratory Research in the Institutes

ESS – European Simulation Symposium

EGG – elliptic grid generation

F

FAS – Full-approximation Scheme

FASTA – Search database that compares a protein sequence to another protein sequence or protein database, or a DNA sequence to another DNA sequence or DNA library.

FBA – flux balance analysis

FEM – finite element method

FETI-DP – Finite Element Tearing and Interconnecting Dual Primal

FFLO – Fulde, Farrell, Larkin and Ovchinnikov, discoverers of ferromagnetic-superconducting state

FFT – Fast Fourier Transform

FIAT – Framework for Interprocedural Analysis and Transformation

FORTRAN – *for*mula *trans*lator, the first compiled high-level programming language.

FOSPACK – a package developed for automatic discretization and solution of FOSLS

FOSLS – First-Order System Least-Squares

FY – fiscal year

G

GAMS – General Algebraic Modeling System

GCC – GNU Compiler Collection

GEM – Geometric Efficient Matching

GHHE – generalized hyperbolic heat equations

GLR – Generalized left right parsers

GMRES – Generalized Minimal Residual (GMRES)

GNU – GNU's not UNIX

GridDB – A Database Overlay for the Scientific Grid

GUI – Graphical User Interface

H

HCI – human-computer interaction

HIV – human immunodeficiency virus

HkDef – Hacker Defender

HLL – Harten–Lax–van Leer, developers of approximate Riemann solver for the Euler equations of inviscid gas dynamics

HP – Hewlett-Packard

HPC – High-Productivity Computing

HTML – HyperText Markup Language

HTTP – Hypertext Transfer Protocol

HVAC – heating, ventilation, and air conditioning

HYDRA – Hydrological routing algorithm that simulates the flow of water

HYPRE – high-performance conditioners

I

IBM – International Business Machines

ICAT – Internet Catalog of Assailable Technologies

ICCD – Integrated Computing & Communications Department

ICCS – Integrated Computer Control System

ICEO – induced-charge electro-osmosis

Acronyms and Abbreviations

ICMT – Internships in Computational Modeling at the Terascale

ICST – Internships in Computer Science at the Terascale

IDS – intrusion detection system

ISO – International Organization for Standardization

IDX – File format used in ViSUS (see) based on a multiresolution space-filling curve index that allows for fast multiscale data access.

ILP – instruction level parallelism

IMAGE – Integrated Molecular Analysis of Genomes and their Expression

IMPACT – Integrated Map and Particle Accelerator Tracking code

IOR – intermediate object representation

IP – Internet Protocol

IPAM – Institute for Pure and Applied Mathematics

I/O – input/output

ISCR – Institute for Scientific Computing Research

ISMG – Information System Management Group

IT – information technology

ITER – International Thermonuculear Experimental Reactor

ITS – Institute for Terascale Simulation

J

JDO – Java Data Objects

JPEG – Joint Photographic Experts Group, a graphics format ideal for photographs, artwork, and paintings; not suited to line drawings, text, or simple cartoons

JSP – Java Server Pages

K

KINSOL – solves nonlinear algebraic systems, see SUNDIALS

KOJAK – POMP-compliant (see) library for profiling and tracing OpenMP (see) applications

KULL – unclassified designation for AX-Division code used to model inertial confinement fusion (ICF, see)

L

LANL – Los Alamos National Laboratory

LDA – mathematical library integer description: On entry, LDA specifies the first dimension of A as declared in the calling (sub) program. Also – local density approximation (first-principles physics approximation within the Linear Expansion in Geometric Objects (LEGO) approach

LDRD – Laboratory Directed Research and Development

LES – Large-Eddy Simulation

LGA – Local Global Alignment

LINPACK– benchmark code for testing supercomputer TF (see) capabilitity by solving systems of linear equations

LLNL – Lawrence Livermore National Laboratory

LOBPCG – locally optimal block preconditioned conjugate gradient

LOCKSS – Lots of Copies Keep Stuff Safe

LOD – level-of-detail

LPI – laser–plasma interaction

LPS – laser proton sources

M

MATLAB – MathWorks proprietary scientific computing and graphics capable programming language, now at V7.0

MCR – Multiprogrammatic Capability Cluster, ALC (see) and MCR are sibling Linux clusters with ~1000 Intel P4 nodes with dual 2.4GHz processors

MDCASK – Molecular dynamics code for radiation damage, to be used as one of the benchmark codes for testing ASC Purple C (see)

MEMS – Micro-Electro-Mechanical Systems

METIS – linear algebra package for partitioning unstructured graphs, partitioning meshes, and computing fill-reducing orderings of sparse matrices, written in FORTRAN

MHD – magneto-hydrodynamics

MIRANDA – research hydrodynamics code ideal for simulating Rayleigh–Taylor and Richtmyer–Meshkov instability growth. Runs on Fortran 95 with MPI (see). Important for these four factors: incompressible and compressible forms; explicit time solution (Poisson solve for incompressible); Eulerian (fixed), Cartesian mesh; and high-order-accurate derivatives

MLIC – multi-layered image cache system

MOMA – minimization of metabolic adjustment

MPI – Message Passing Interface.

MRF – Markov Random Field

fMRI – functional magnetic resonance imagingMRO – Mars Reconnaissance Orbiter

N

NAI – Nonproliferation, Arms Control, & International Security Directorate at LLNL

NAIC – NAI (see) and Computing Applications Division within CAR (see)

NAMD – object-oriented molecular dynamics code designed for simulation of large biomolecular systems

NASA – National Aeronautics and Space Administration

NCAR – National Center for Atmospheric Research located in Boulder, CO; and the software products NCAR Command Language (NCL) and NCAR Graphics that facilitate forecasting and visualization

NFS – Network File System

NIF – National Ignition Facility

NIFE – NIF (see) and Engineering Computing Applications Division within CAR (see)

NNSA – National Nuclear Security Administration

NR – non-relativistic

O

ODE – ordinary differential equation

OODT – object-oriented data technology

OpenMP – open message-passing, de facto standard for shared-memory parallel programming of scientific applications

ORNL – Oak Ridge National Laboratory

OS – operating system

OSVDB – Open Source Vulnerability Database, a project to catalog and describe global security vulnerabilities, opened a vendor dictionary as a centralized resource on August 31, 2004. OSVDB is sponsored by Digital Defense, Inc.(1999), a private global network security provider; and by Churchill & Harriman (1986), security business partner to mid-market and Fortune 500 companies, headquartered in Princeton, NJ.

P

PAPI – Performance Analysis Programmable Interface sfotware tool: open, cross-platform interface to the performance analysis hardware found in most modern microprocessors

PAT – Physics and Advanced Technologies Directorate at LLNL

PaToH – Partitioning Tools for Hypergraph

PC – personal computer

PCG – preconditioned conjugate gradient

PCMDI – Program for Climate Model Diagnosis and Intercomparison

PCR – polymerase chain reaction assays can amplify a target segment of DNA in suspect biological organisms

vPCR – Virtual PCR (see above)

PDB – Protein Data Bank

PDEs – partial differential equations

PDS – Planetary Data System

PERC – Performance Evaluation Research Center, a SciDAC (see) integrated software infrastructure center with four strategies for maximizing memory hierarchy: discipline-specific benchmarks; performance analysis tools; performance modeling; and performance optimization tools

PerfTrack – a database-based tool for storing, navigating, and analyzing very large amounts of performance data

PF3D – LLNL 3D laser–plasma interaction code

PHP – recursive acronym for "PHP: Hypertext Preprocessor": general-purpose scripting language well suited for Web development and easy to embedded into HTML, commonly used with the Apache HTTP server and included in Red Hat Linux versions

PHP/MySQL – These functions allow the user to access MySQL database servers. More information about MySQL can be found at http://www.mysql.com/

PI – Principal Investigator

PIMS – the LLNL Engineering Directorate Personnel Information Management System

PIW – Promoter Identification Workflow

PLASMA – PLAnetary Scale Monitoring Architecture

PMaC – Performance Modeling and Characterization

PMPI – performance-monitoring programmable interface

POMP –standard performance monitoring interface for OpenMP (see), an API (see) to be called by probes inserted into the application by a compiler, a pre-processor, or via a binary or dynamic instrumentation mechanism

POP – Parallel Ocean Program

POSIX – Portable Operating System Interface incorporates the IEEE and Open Group set of fundamental services needed for the efficient construction of application programs.

PROLINKS – database for co-evolving proteins, used in biological studies and comparisons

PUM – partition of unity methods

PVODE – parallel ODE (see) integrator, a special case of the scaled nonlinear solver, see SUNDIALS

Acronyms and Abbreviations

Q

QBF – query by form

QC – quality control

QCD – Quantum Chromodynamics

R

RANS – Reynolds-Averaged Navier–Stokes equations

RANSAC – Random Sample Consensus

RBF – radial-basis functions

REB – relativistic electron beams

ReiserFS – journalling filesystem included in Linux 2.4, designed and developed by Hans Reiser and his team at Namesys at, creates a single shared environment, or namespace, where applications can interact more directly, efficiently and powerfully. Users can access the filesystem directly rather than building special-purpose layers that run on top of the filesystem, such as databases.

RKPM – reproducing kernel-particle methods

RM – Richtmyer-Meshkov

RMI – Remote Method Invocation.: Java RMI allows the user to invoke a method on an object that exists in another address space — on the same machine or a different one.

RNA – ribonucleic acid: Genetic code is stored in the DNA sequence, which is transcribed into RNA and translated into a polypeptide—proteins, enzymes, or peptide hormones.

ROAM – Real-time Optimally Adapting Meshes: AMR (see) tool in which two priority queues drive split and merge operations that maintain continuous triangulations built from pre-processed bintree triangles

ROSE – name, not an acronym for the Overture Suite preprocessor that recognizes user-defined objects and substitutes optimized code before compilation

RTRAM – Real-Time Radiation Area Monitoring Network

RTS – runtime system

S

SAGE IR – at the request of LANL (see), the SAGE ASC (see) benchmark code is no longer publicly accessible

SAMRAI – Structured Adaptive Mesh Refinement Application Infrastructure

SCaLeS – Science-based Case for Large-scale Simulation

SCD – service-class description

SCE – statistical condition estimation

SCI – Scientific Computing and Imaging Institute at the University of Utah utilizes component-based environments for biomedical computing, computational combustion and other applications

SciDAC – Scientific Discovery through Advanced Computing

SCons – Open Source Software Carpentry tool that is, a next-generation build tool, an improved, cross-platform substitute for the classic `Make` utility with integrated functionality similar to `autoconf/automake` and compiler caches such as `ccache`.

SCWRL – Side Chain placement With a Rotamer Library

SDM – scientific data management

SDSC – San Diego Supercomputer Center

SEGRF – Student-Employee Graduate Fellowship

SGS – Slow Growing Subdivision

SIAM – Society for Industrial and Applied Mathematics

SIDL – Scientific Interface Definition Language

ccSIM – cache-coherent memory simulator

SLURM – Simple Linux Utility for Resource Mnaagement: Open Source, fault-tolerant, and highly scalable cluster management and job scheduling system for large and small Linux clusters

SMPs – symmetric multiprocessors

SPASE – Space Physics Archive Search and Extract

SOAP – an emerging communication standard that encodes remote method invocations using XML payloads over network transport mechanisms such as HTTP

SPH – smoothed-particle hydrodynamics

SPMD – Single Program Multiple Data, parallel programs that use multiple processes running the same code working on different data to solve a problem

SQA – software quality assurance

STL – Standard Template Library

STREAM – Stanford Stream Data Manager

SK – SucKIT – a fully working rootkit that is loaded through `/dev/kmem`

SUIF – Stanford University Intermediate Format

SUNDIALS – (**SU**ite of **N**onlinear and **DI**fferential/**AL**gebraic equation **S**olvers) consists of the following four solvers.

CVODE solves initial value problems for ordinary differential equation (ODE) systems.
CVODES solves ODE systems and includes sensitivity analysis capabilities (forward and adjoint).
IDA solves initial value problems for differential-algebraic equation (DAE) systems.
KINSOL solves nonlinear algebraic systems.

SUPRI – Stanford University Petroleum Research Institute, research group interested in the design of efficient and accurate simulation tools for compositional problems, such as those occurring in gas injection processes

SVM – Support Vector Machine

SWA – segmentation by weighted aggregation

T

TAMM – Terrestrial and Atmospheric Monitoring and Modeling

TAU – Tuning and Analysis Utilities, a program and performance analysis tool framework for high-performance parallel and distributed computing

TCP – transmission control protocol

Teraflop/s or TF – trillion floating-point operations per second

TIFF – Tagged Image File Format, a file format used for scanning, storage, and interchange of gray-scale graphic images

TPS – thin-plate spline

TRANSFAC – Transcription Factor, BIOBASE proprietary database on eukaryotic transcription factors, their genomic binding sites and DNA-binding profiles.

TSTT – Terascale Simulation Tools and Technologies

U

UC – University of California

UCI – University of California, Irvine

UCRP – University Collaborative Research Program

UI – user interface

UIUC – University of Illinois, Urbana-Champaign

URL – Uniform Resource Locator

URP – University Relations Program

UWB – ultra-wideband

V

VBL – Virtual Beam Line

VETFEM – Variable-Element-Topology Finite Element Method, a general-purpose finite in which each element is free to take essentially any polygonal (polyhedral in 3D) shapeVisIT – contraction of Visualize It, a free DOE/ASC (see both) interactive parallel visualization and graphical analysis tool for viewing scientific data on Unix and PC platforms

ViSUS – Visualization Streams for Ultimate Scalability

VTDB – Vulnerability Tracking Database

VTK – Visualization Toolkit

W

WHIRL – Word-based Information Representation Language, a representation system that combines some of the properties of relational databases, and some of the properties of statistical ranked-retrieval systems.

WPI – Worcester Polytechnic Institute

X

XML – eXtensible Markup Language

XWRAP – an XML-enabled software system for semi-automatic generation of wrapper programs for Web sources

Summer Research Program

The ISCR put on its largest summer visitor program ever in FY 2004 with approximately 100 students in residence, as well as 11 faculty. Students are assigned individual LLNL mentors and given specific projects, ranging from programming tasks to original research, to which they will contribute based on their background and skills.

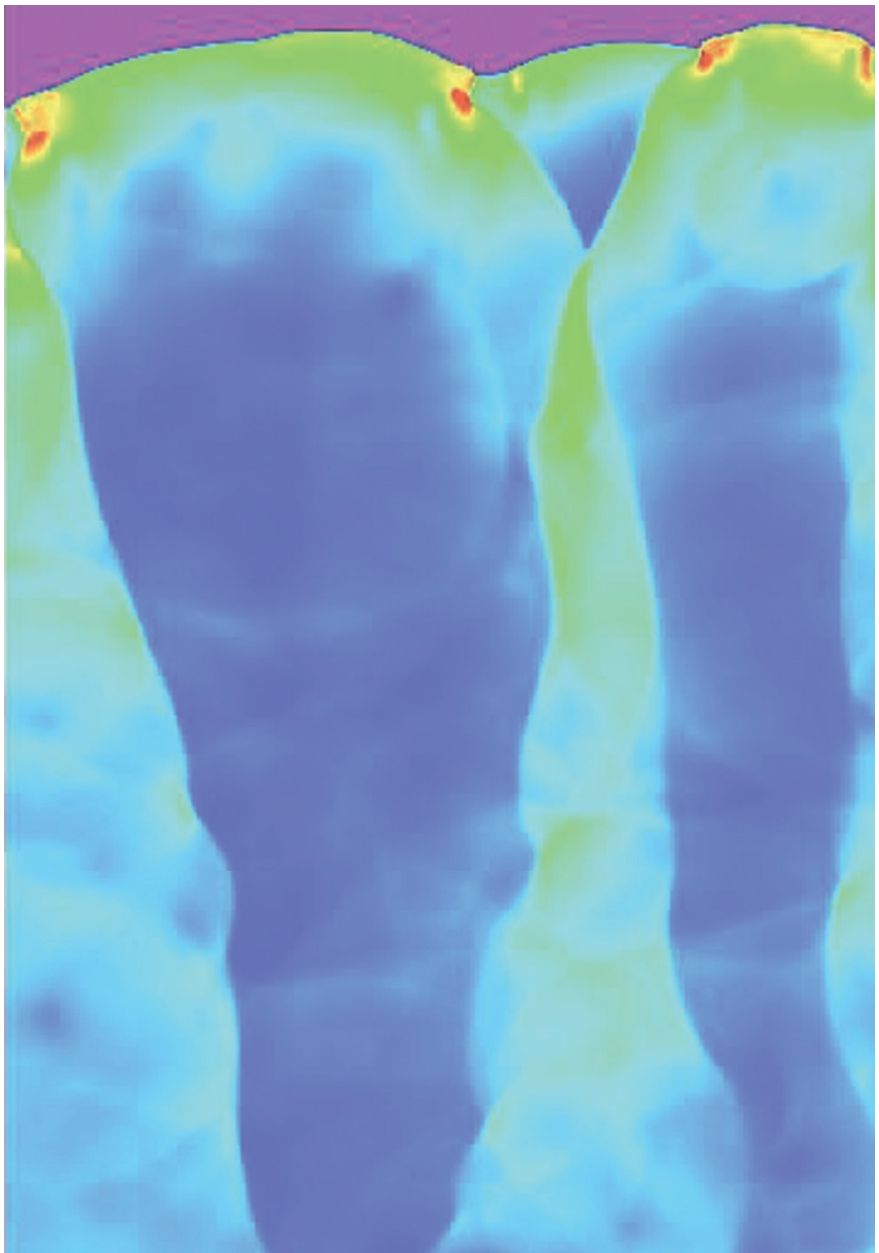

TOC

Unstructured Mesh Technology Development

Student

Carl F. Ollivier-Gooch, University of British Columbia

Mentor

Lori Diachin, CASC

The research proposed under this subcontract consisted of three parts:

1. Enable Mesquite to use GRUMIVW topology modification routines and enable GRUMMP to use Mesquite's new vertex smoothing technology
2. Explore possibilities for globally optimal mesh reconnection.
3. Attempt to create a high-order accurate unstructured mesh version of a computational electromagnetics code, and explore the possibilities for adding high-order accurate unstructured mesh capabilities to Overture.

This report describes the work done on each of these three topics and future prospects in each area.

1. Mesquite-GRUMMP Interconnection

This task decomposes naturally into two sub-tasks: enabling communication of mesh data between two large pieces of unstructured mesh manipulation code that use very different data structures, and supplying the small amount of control information to specify what the mesh improvement goal is.

Communication of mesh data between GRUMMP and Mesquite is simplified, or at least given a specific structure, by the Terascale Simulation Tools and Technologies (TSTT) mesh interface definition. While at LLNL, I implemented the TSTT mesh interface using the GRUMMP mesh database as the back end. In the process, I also identified several missing or ambiguous items in the interface and made proposals to the TSTT interface group to correct them. Major issues that I have helped to clarify in the TSTT interface include error handling and entity sets. At this point, the GRUMMP implementation of the TSTT interface, while not yet 100% complete, is more than sufficient to support all operations required by Mesquite for vertex smoothing.

The reverse communication – Mesquite implementation of the TSTT interface – is not yet complete.[1] By the time it is, implementation of mesh reconnection using the algorithms currently in GRUMMP but with mesh changes through the TSTT interface should also be complete. This will enable not only mesh reconnection in Mesquite but in other TSTT-based tools, as well. Work on the TSTT mesh reconnection service has already begun.

Once full two-way mesh data communication between Mesquite and GRUMMP is possible, implementing GRUMMP's mesh reconnection in Mesquite and Mesquite's new vertex smoothing in GRUMMP will require little more than specifying the quality measure to optimize and calling the optimization routine.

2. Global Mesh Reconnection

The goal of this part of the project was to determine the feasibility of using global rather than local mesh reconnection techniques to improve unstructured mesh topology. I began by creating a naive implementation of the well-known branch-and-bound strategy for combinatorial optimization. The initial implementation demonstrated that global reconnection is possible, but execution times were absurdly high. A series of improvements to the algorithm were made, all aimed at improving the effectiveness of branch pruning, because effective pruning can dramatically reduce the number of possible mesh configurations that must be examined in determining the best global connectivity. As a result of these efforts, the original code was sped up by a factor of over 1000 for small meshes. Unfortunately, run times are still relatively slow, and the trend in run time with mesh size is poor: in three dimensions, run time scales roughly with the seventh power of mesh size, which is clearly prohibitive. Nevertheless, it was felt that the ideas involved were of sufficient interest to the meshing community to justify a conference paper, and a draft paper was submitted to the 13th International Meshing Roundtable. (Attached)

[1] Mesquite coding will be done by a DOE employee or post-doc, so timing on that is out of my control.

Carl F. Ollivier-Gooch, (Continued)

3. High-order Discretization

The main effort here was in discussions with the Overture team about the possibility of adding high-order unstructured operators to the Overture framework. I learned how Overture is constructed and how it provides various differential operators efficiently at high orders of accuracy for structured meshes. I also discussed at some length with Bill Henshaw and Kyle Chand of the Overture team how one might implement these types of operators using the high-order unstructured mesh solution reconstruction techniques that my solver library provides. While we all agreed that using my high-order reconstruction code to extend Overture's capabilities is both useful and feasible, we reached no firm conclusions about how and when to proceed with implementation.

Boundary Integral Equation Methods and Software for Poisson's Equations in Three Dimensions

Student

Anne Greenbaum, Professor of Mathematics, University of Washington

Eric Machorro, Graduate Student, University of Washington

Mentor

Britton Chang, CASC

During my sabbatical quarter at LLNL, I worked on two projects. One involved the fast solution of Poisson's equation on arbitrary 3D regions, and the other involved finding new methods and analyzing current methods for solving the linearized Boltzmann equation for neutron transport.

For the first project, I implemented a code for solving Poisson's equation on an arbitrary 3D region by first embedding the region in a cube. Next, I solved a boundary-integral equation to determine the right-hand side vector for a fast Poisson solver and then applied the fast Poisson solver on the cube. This algorithm had been described in the literature (1) but had not previously been implemented as a package. In doing so, I unfortunately discovered a few drawbacks.

First, the boundary integral equation must be solved to greater than second-order accuracy in order to obtain second-order accuracy in the Poisson solver. This can be done using a form of Richardson extrapolation, but it means that the user must supply more than a simple triangulation of the surface. Next, while the boundary integral equation can be solved with just a few GMRES iterations, the cost of forming the matrix and multiplying the matrix at each iteration is great. To make the method competitive, these matrix-vector multiplications need to be carried out with the fast multiple method. Finally, an improved procedure for calculating boundary values on the cube needs to be implemented. The conclusion was that the current code is not yet competitive with other methods, such as multigrid methods, for solving Poisson's equation on interior regions, although it might be the method of choice for exterior problems where finite difference methods may not be applicable.

The other project I worked on, along with graduate student Eric Machorro, had to do with neutron transport. We were particularly interested in methods that would avoid negative fluxes and oscillations, yet would maintain second-order accuracy or better. Eric implemented the exponential characteristic (EC) method (2), which avoids negative fluxes and maintains conservation but is considerably more expensive per cell than the more standard diamond-difference or Petrov-Galerkin methods. Moreover, we were able to show that it is only first-order accurate in the mathematical sense, although it may obtain high accuracy for certain problems anyway. Working with Britton Chang of LLNL, we were able to rigorously establish second-order accuracy for the Petrov-Galerkin method, a property that had previously been suggested but never proved. A paper on this has been written and will soon be submitted for publication.

References

1. Greenbaum and Mayo, "Rapid Parallel Evaluation of Integrals in Potential Theory on General Three-Dimensional Regions," *Jour. Comp. Phys.*, 145 (1998), pp. 731-742

2. Minor and Mathews, ``Exponential Characteristic Spatial Quadrature for Discrete Ordinates Radiation Transport with Rectangular Cells," *Nuclear Sci. Engr.*, 120 (1995), pp. 165-186

Embedded Boundary Methods for Partial Differential Equations

Student

Bjorn Sjögreen, Royal Institute of Technology, Stockholm, Sweden

Mentor

Anders Petersson, CAR, CASC

First, we investigated far-field boundary conditions for Maxwell's equations. After a preliminary study of different types of boundary conditions, we decided to use the perfectly matching layer (PML) technique. Although this is a well-known method, there are several parameters in the equations that are usually selected based on computational experience. In my work, I have been able to show mathematically exactly how these parameters should be chosen in order to have a grid convergent method. This far-field boundary condition was implemented to compute scattering from perfectly conducting objects. The objects were represented as embedded curves in a Cartesian grid using the new embedded boundary method by H.O. Kreiss and A. Petersson. As a second project, I redefined the above embedded boundary method for Maxwell's equations to be able to solve compressible fluid flow problems.

The original embedded boundary method did not perform well for such problems because of the shockwaves that usually exist in compressible fluids. It was therefore necessary to modify the embedded boundary interpolation formulas. I introduced special limiter functions into the boundary interpolation and narrowed the interpolation stencil. With these changes, the method performs very well. An example of flow past three disks at Mach 3 is shown below. Finally, I have implemented the method for three space dimensions and computed compressible fluid flow past a sphere.

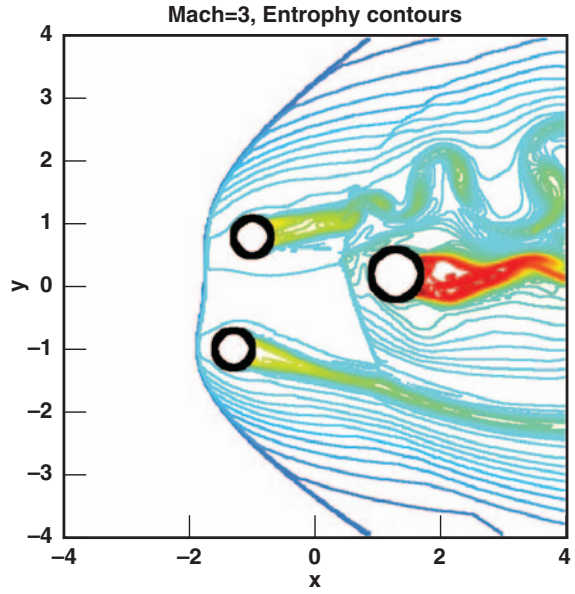

The figure shows streamlines for air flowing in from the left at supersonic speed. A strong bow shock is formed in front of the cylinders, a smaller shock can be seen in front of the rearmost cylinder.

TOC

A New Approach for Solving the Stokes Problem Based on the Distributed Relaxation Method

Student
Constantin Bacuta, University of Delaware

Mentor
Rob Falgout, CASC
Panayot S. Vassilevski, CASC

We present a distributed relaxation method for the incompressible Stokes problem. The method can be applied to discretizations of the Stokes system provided that the discrete approximation spaces for velocity and for pressure form a stable pair. We establish sharp rates of convergence in terms of the inf-sup condition constants at the continuous and the discrete level. We prove that our iterative process is optimal among all Uzawa-type algorithms that involve constant step-size parameters. Based on the functional iteration formulation of the new algorithm, we are able to analyze Achi Brandt's distributed relaxation method as an iterative process and compare it with Uzawa-type methods for solving the Stokes system. The distributed relaxation method for the incompressible Stokes problem is based on a change of variables that leads to a lower triangular system with Laplace operators on the main diagonal for which multigrid methods are more suitable.

We propose a finite-element formulation based on the distributed relaxation method. Using the new approach, we can construct effective and robust multigrid methods for solving Stokes-type systems and other PDE elliptic systems that can be reformulated as saddle-point problems.

We plan to present numerical experiments to demonstrate the effectiveness of the transformation, which is well-established for certain finite difference discretizations of Stokes problems.

For a future research activity, we consider the Maxwell equations, reformulated as saddle point problem. With an appropriate change of variables, a similar iterative process can be used to obtain an efficient way to discretize and solve the Maxwell equations.

References

Bramble, J.H., J.E. Pasciak and P.S. Vassilevski, "Computational scales of Sobolev norms with application to preconditioning," *Mathematics of Computation*, 69 (2000), pp. 463-480.

Brandt, A., and N. Dinar, "Multigrid solutions to flow problems," *Numerical Methods for Partial Differential Equations,(S. Parter, ed.),* 1979, Academic Press, New York, pp. 53-147.

Girault, V. and P.A.Raviart, "Finite Element Methods for Navier-Stokes Equations,"
Springer-Verlag, Berlin, 1986.

Wittum, G., "On the convergence of multi-grid methods with transformed smoothers,"
Numerische Mathematik, 57 (1990), pp. 15-38.

TOC

Compatible Relaxation and Optimal AMG Interpolation

Student

James Brannick, University of Colorado, Boulder

Ludmil Zikatanov, Pennsylvania State University

Mentor

Rob Falgout, CASC

Panayot Vassivlevski, CASC

In collaboration with Rob Falgout and Panayot Vassivlevski, we worked on the development and implementation of some state-of-the-art techniques for the construction of Algebraic Multigrid Methods (AMG) for the solution of positive definite linear systems arising from the discretizations of elliptic partial differential equations. Our approach uses the fact that a Compatible Relaxation (CR) process reduces the condition number of the matrix corresponding to the fine grid (usually denoted with Aff) as much as possible. Our idea is that the supports of the coarse grid basis functions (Ωi) can be determined by constructing an approximation to Aff-1Afc (Afc represents the coupling between coarse-grid and fine-grid degrees of freedom). Once the supports are defined, we determine the basis in a way that the trace of the coarse grid matrix is as small as possible.

Such a scheme can be made quite robust, by dynamically taking better and better approximation to Aff-1Afc. This is done by increasing the size of Ωi or updating the approximation to the slow-to-converge error, which is obtained by applying the iteration to the homogeneous system of equations Au = 0 and then and recomputing the basis during the V-cycle iterations. Our current studies show that a well-conditioned Aff, provided by a CR procedure, leads to a good approximation, even by choosing Ωi to be a small size.

The set of the test problems we have used so far is on matrices corresponding to discretizations of scalar PDEs. Our plan is to extend these techniques and implement them in the case of vector equations and systems of PDEs, such as linear elasticity and Maxwell equations. Although it may not be optimal, we are confident that an ad hoc application of the idea described above will work in these more complicated cases.

TOC

Mixed and Multigrid Methods for Stokes and Navier-Stokes Equations

Student

Zhiqiang Cai, Purdue University

Mentor

Panayot Vassilevski, CASC

To discretize the incompressible Stokes and Navier-Stokes equations by the finite element method, one requires a stable combination of finite element spaces in order to approximate the velocity and pressure in a well-defined and optimal manner. Commonly used stable pairs include the Taylor-Hood element. The algebraic equations resulting from the discretization is symmetric but indefinite. Many solution methods that work well for symmetric and positive definite problems cannot be applied directly. Although substantial progress in solution methods for saddle-point problems has been achieved, these problems may still be difficult and expensive to solve.

The purpose of this project is to develop, analyze and implement accurate and efficient numerical methods for solving the Stokes and Navier-Stokes equations. Instead of using the velocity-pressure formulation, our method is based on the pseudostress-velocity formulation developed jointly with LLNL scientists. The pseudostress is approximated by the Raviart-Thomas (RT) and the velocity by piecewise discontinuous polynomials. We have shown that this pair of finite elements is stable and yields quasi-optimal accuracy. Important physical quantities, such as the pressure, vorticity and stress, may be calculated algebraically through post-processing.

To efficiently solve the indefinite system of linear equations resulting from the discretization of the stationary Stokes equation, we decouple the system by the penalty method. The penalty parameter is chosen to be proportional to the approximation accuracy of the discretization for avoiding accuracy loss. Even though the reduced pseudostress system is very ill-conditioned, we have developed an efficient multigrid method based on the H(div)-type of multigrid method. The velocity can then be recovered without solving any equations. Numerical experiments on uniform rectangular RT elements of the lowest order show that multigrid convergence factor is around 0.25 and that it is independent of the mesh size, the number of levels, and the penalty parameter.

In the future, we will study efficient multigrid methods for solving system of linear equations resulting from the discretization of the time-dependent Stokes equations. This is very challenging because the resulting pseudostress system at each time step has many extra near-null space components. We will also continue our study for both stationary and time-dependent Navier-Stokes equations.

TOC

Adaptive Multigrid via Subcycling on Complementary Grids

Student **Timothy Chartier,** Davidson College

Mentor **Edmond Chow,** CASC

Algebraic multigrid methods uncover or assume the nature of algebraically smooth error components in order to construct multigrid components. Adaptive multigrid schemes expose algebraically smooth error, analyze the effectiveness of the resulting multigrid algorithm and adjust the cycling as needed in order to improve the rate toward convergence.

This research, conducted in collaboration with Edmond Chow, focused on the use of relaxation and subcycling on complementary grids as an evaluative tool in correcting multigrid cycling. Each complementary grid is constructed with the intent of dampening a component of algebraically smooth error. In particular, complementary grids are constructed when the adaptive cycle determines that smooth modes are not being sufficiently eliminated by the current multigrid process. The particular implementation of this idea manages smooth error in a manner analogous to spectral AMGe.

The algorithm was tested on a model problem with Dirichlet boundary conditions. On both the isotropic and anisotropic Laplacian operators, convergence rates for 2-4 levels remained at less than 0.10. It is important to note that comparable results occurred when the system $Ax = b$, was scaled to produce the new system $SASx = b$

where S was a diagonal matrix with $s_{ii} = 10^{d_i}$ and $0 < d_i < 103$. For the rotated anisotropic diffusion problem, convergence rates were kept below 0.20 for 2-4 levels. Again, the scaling, which troubles traditional algebraic multigrid, produced comparable results.

Future numerical tests will include larger systems and a more extensive pool of problems.

TOC

AMG Algorithm for Finding an Eigenbasis for the Schroedinger Operator

Student

Irene Livshits, University of Central Arkansas

Mentor

Robert Falgout, CASC

The goal of this project is to develop multigrid solvers that allow an efficient approximation to eigenvalues and eigenfunctions of Schroedinger operators with variable potentials. The developed method allows fast and inexpensive calculation of all eigenvalues in a given interval, as well as the corresponding eigenfunctions. The calculated eigenbasis has a multilevel structure – the number of eigenfunctions approximated on each grid grows as grids become coarser. All eigenfunctions are approximated and the eigenvalues are calculated only on the coarsest grid. This structure is not only inexpensive to calculate and store, it also allows for efficient use of the eigenbasis, e.g. for a fast summation.

The algorithm deals with the indefinite nature of a Schroedinger operator using the wave-ray approach developed by Livshits and Brandt for Helmholtz-boundary value problems with constant potentials. However, unlike the geometric wave-ray method, the new algorithm employs algebraic multi-Galerkin approach to construct coarse grid operators. To date, the solvers were developed for one-dimensional operators, serving as a calibration for two- and three-dimensional algorithms.

The suggested approach not only aims to solve computationally challenging Schroedinger equations, it can be also used for other operators that have several independent low eigenfunctions.

TOC

Large-scale Integration of Web Sources

Student
Anne Ngu, Texas State University

Mentor
Terence Critchlow, CASC

I worked on two specific aspects of this project during this summer. The first aspect is a continuation of last summer's work in the automatic discovery and classification of Web sources. My task is to apply the metadata-based classification methodology used for identification of BLAST servers to other application domains. This involves creating a service class description for bio-keyword search applications and conduct experiments to measure the performance of our classifier for this domain. The result is incorporated in the paper submitted to the World Wide Web Journal.

The second aspect of my work this summer was to investigate the possibility of automatically generating a service class description for automatic classification. The manual creation of service class description is tedious, error prone and not adaptative to the dynamic nature of Web sources. The idea is to use a supervised learning technique for automatically generating the regular expressions that can be used for recognizing the data embedded in an HTML response page. Given training documents and a set of user-annotated examples, including a few sample sites of a class of Web source, the system will automatically generate a set of data type rules in the form of regular expressions that can be used for the correct identification of all the sources that belong to that class.

I implemented a clustering-based learner algorithm for automatic generation of data type rules and incorporated it with the existing dataRecognizer system at LLNL. The initial experimental evaluation of the generated rules showed them to be superior with respect to the quality of rules and document classification time. A report entitled "Automatic Generation of Data Types for Classification of Deep Web Sources" is being prepared.

Adapting Algebraic Multigrid for the Solution of the Curl-Curl Formulation of Maxwell's Equations

Student **David Alber,** University of Illinois at Urbana-Champaign

Mentor **Barry Lee,** CASC

Applying multigrid to solve linear systems arising from the discretization of Maxwell's equations has been a difficult problem largely because of the large null space of the curl-curl operator. Other solvers, such as Krylov subspace solvers, will work on such systems. However, the desire to have a solver that will scale effectively on a larger numbers of processors is a great motivation to adapt multigrid to the task. With more than 65,000 processors, BlueGene/L is a good architecture to highlight the need for scalable solvers.

Multigrid methods have been created to solve Maxwell's problems, but these methods tend to degrade in performance when the values of the coefficients in the problem vary. Jim Jones and Barry Lee recently introduced a method designed to handle these more complicated problems effectively by allowing the curl- and divergence-free near-null space error components to be annihilated separately. The method shown in their paper is for use on two-dimensional problems on structured grids. Its viability on three-dimensional structured problems was shown by Dylan Copeland.

The purpose of this work was to begin creating a solver based on the same principles that work in two-dimensions on problems defined on unstructured grids. This summer, tests were run in an attempt to approach this problem in different ways using an algebraic multigrid. It is clear that more sophisticated ways of addressing the problem using algebraic multigrid need to be investigated. In particular, an effective way to select coarse grids for curl- and divergence-free errors needs to be formulated.

Locally Optimal Methods to Solve Eigenvalue Problems in Electronic Structure Calculations

Student **Kris Andersen,** University of California, Davis

Mentor **John Pask,** H Division: Metals and Alloys Group

The goal of my project was to find the most efficient way to implement a new locally optimal method to solve generalized eigenvalue problems – developed by Andrew Knyazev and Richard Lehoucq within the context of electronic structure calculations. Since the matrices involved are large (106) and the tolerances are more stringent than typical engineering applications, this involved paying close attention to compromises between efficiency and robustness. Specifically, I developed a new way to handle deflation that makes use of the reasonable eigenvector approximations available in an electronic structure calculation. I performed a number of numerical experiments to find the most efficient way to do specific operations, such as restarting the conjugate gradient algorithm. I also looked closely at preconditioning and learned how to interface our code (which uses complex arithmetic) with the HYPRE package (which only supports real numbers). I also explored how to use numerically cheaper preconditioners, such as a PCG linear solver, more effectively.

TOC

QTester

Student **John Anderson,** University of California, Davis

Mentor **Benjy Grover,** DCOM

This summer, I rearchitected the QTester application, a graphical user interface testing tool. QTester makes it easy to design and run regression tests via the graphical user interface of any Qt-based application. Users can create test cases while using their target application. The QTester works by recording Qt events, such as mouse clicks and key presses, and then replaying the events upon request. For most applications, the QTester opens up greater opportunities for assuring software quality.

Real-Time Radiation Area Monitoring: Emergency Response and Regulation

Student **Dustin Anderson,** California Polytechnic State University, San Luis Obispo

Mentor **Don MacQueen,** Environmental Protection Department

Sixteen Geiger-Mueller sensors set up around the Lawrence Livermore National Laboratory continuously monitor for harmful radiological conditions. Using meteorological data from the Lab's meteorological tower, the Real-Time Radiation Area Monitoring Network (RTRAM) compares current radiation counts to reference values to determine threat severity. The networked sensors on site are connected to a Central Command Center where their radiation readings are rerouted in real-time. Incorporated into this network is an electronic notification system that immediately alarms the Emergency Operations Center about hazardous conditions. If a situation arises, the threat is automatically reported to the EOC so that appropriate measures may be taken in a timely manner. The RTRAM network's purpose at the Lab is twofold; it doubles as both a safety net for the Lab and ensures that daily radiation emissions are within regulations.

The software managing RTRAM's web site is written in Perl, SQL, and now JavaScript. My participation in the project included working on these scripts to optimize how data is retrieved from the database, how the data is organized and formatted in the real-time viewable web pages, and how the data is stored, adding functionality, and fixing compatibility issues while converting old data to new data. This new means of storing data improves real-time performance, which increases usability and improves the efficiency of the entire RTRAM network. Boosting performance allows users to see data faster than ever and helps them respond to critical situations more quickly.

Related Projects

In order to be better prepared to renovate the RTRAM network and make my work on it as efficient as possible, I took on several side projects to help me better understand the programming languages Perl, JavaScript and SQL, as well as the nuances of HTML and XML. The other projects I worked on as pseudo-training before tackling the RTRAM network software included, but are not limited to:

Oracle Password Pages

The first project that introduced me to the Perl programming language and also to database communication using Perl and SQL was a set of two pages that enabled users to change their password on the Oracle database. Users would enter all relevant information and after error-checking all entered data, the user's password would be changed as long as all tests were passed. This was also my first exposure to SQL, a language designed for database communication through the use of queries.

QC Chemist Validation Web Pages

These Quality Control (QC) Web pages were written as a replacement for a tool on the old Ingres server used by the ERD ISMG chemists. Users can now access all chain of custodies (COCs) on the Oracle database by entering a COC ID and then altering and adding data as necessary. These pages, and the DMT Receive Data Web Pages described below, were critical milestones in ERD ISMG's target goal of going parallel and switching from the Ingres database system to the Oracle database system.

Data Management Tools (DMT) Receive Data Web Pages

Similar to the QC Chemist Validation pages, the DMT Receive Data pages allow data managers to access different aspects of all chain of custodies on the Oracle Database by entering a desired COC ID. Information regarding a specific COC can be altered and/or deleted through these pages, allowing data managers to both double-check and enter new queries.

CES Algorithm Review Web Pages

The Cost Effective Sampling (CES) Algorithm Review Web pages are used by the Environmental Restoration division of the Environmental Protection Department to help determine how often it samples ground water monitoring wells. Altering these pages included adding automatic email verification and redirection of pages, as well as formatting the pages

Dustin Anderson (Continued)

to be consistent with the other data management tools (DMTs).

Drop Tables Pages

Enhancing the previous "Drop Tables" Web page, the new Drop Tables pages allow data managers to drop tables from the Oracle database. Previously, users had no choices to make, but the new pages allow users to drop either individual tables or all tables at once. This was my first encounter with JavaScript, and I used these skills later when I altered the RTRAM network scripts.

The Terrestrial and Atmospheric Monitoring and Modeling Group Web Page

This was my main training in using HTML and XML to format data. Creating a web page with the use of explicit table layouts, the TAMM group's overview web page is now both informative and aesthetically pleasing.

Documentation of the RTRAM Web Site

In order for others who join the RTRAM project to have an easier time learning, understanding and navigating the network, I created 20 separate Web pages written in HTML/XML and linked together as an "Electronic Handbook." The Handbook can be browsed as a normal Web site, and each page summarizes exactly one Perl script that runs the site with links mimicking those on the actual site. This document provides summaries of each of the Perl scripts, along with showing how each script relates to all the others.

Query By Form (QBF) Tool: Group Delete Functionality

Adding functionality to the already existing QBF tool required me to add elements to the Web tool in order to make the data managers' jobs a little easier. Basing my code on the Drop Tables pages mentioned above, I enabled group delete functionality to this tool, giving data managers the ability to delete specific rows from their data tables as individuals or as user-defined groups.

Adding New Functionality to the RTRAM Website

Although many features were altered and/or optimized during "renovation," two significant new features and four Perl scripts were added. Comments can now be added to censor data during specific times and/or time intervals. These comments can also be easily viewed through the Web site – features that were not previously offered.

Converting Old Data & Compatibility

Old data, stored in unusable format, can now be converted to the new data format easily with a Perl script. Any number of data files in the old format can be converted to the new format simultaneously, enabling data retrieval via the Web to produce a *complete* data set, rather than the partial set that was previously created.

PIMS Regression Test Development and Product Enhancement

Student **Benjamin Apodaca,** Northern Arizona University

Mentor **Carolyn Wimple,** NIFE

This summer I worked on three main tasks: creating testing scripts to perform regression testing on a space utilization system called Engineering Facility Information (EFI), creating a prototype of the Skills section of the Personnel Information Management System (PIMS), and creating a Web version of the EFI User Instruction document.

I was introduced to a software tool called WinRunner, which I had never used before, and used it to create the test scripts for EFI. I was also introduced to HTML and JavaScript, and I used these Web programming languages to create both the Skills prototype and the Web version of the EFI User Instruction document. All three of the tasks are part of a larger project focused on the PIMS, which is still being performed by the team on which I worked.

Design and Implementation of an Anomaly Detector

Student **Abraham Bagherjeiran,** University of Houston

Mentor **Erick Cantú-Paz,** CASC

I designed and implemented a general-purpose anomaly detector for streaming data. Based on a survey of similar work from the literature, a basic anomaly detector builds a model on normal data, compares this model to incoming data, and uses a threshold to determine when the incoming data represents an anomaly. Models compactly represent the data but still allow for effective comparison.

Comparison methods determine the distance between two models of data or the distance between a model and a point. Threshold selection is a largely neglected problem in the literature, but the current implementation includes two methods to estimate thresholds from normal data.

With these components, a user can construct a variety of anomaly detection schemes. The implementation contains several methods from the literature. Three separate experiments tested the performance of the components on two well-known and one completely artificial dataset. The results indicate that the implementation works and can reproduce results from previous experiments. For more details, see the technical report on the Sapphire Web site at http://www.llnl.gov/casc/sapphire/

I also update and repaired problems with the existing Sapphire code base.

Scalable Graph Algorithms

Student **Paul Baginski,** University of California, Berkeley

Mentor **Van Henson,** CASC

My project involved splitting a graph into topological neighborhoods of a fixed radius. In other words, given a graph G and a natural number d, pick a subset S of the vertices such that every vertex of G is within distance d of some element of S. The vertices in S are called centers, and for a given center, the set of all vertices of G within distance d of that center is called the d-neighborhood. Our goal was to find a relatively small set of centers.

Three distinct algorithms were produced—the Degree algorithm, Tree algorithm, and Complement algorithm. We measured the effectiveness of these algorithms using the following criteria:

(a) Relative number of centers chosen
(b) Runtime
(c) Scalability of the algorithm to different values of d
(d) Amount of overlap in the d-neighborhoods
(e) Amount of information needed beforehand (e.g., adjacency matrix).

Both the Degree algorithm and the Tree algorithm performed very well under most criteria, but in each case there was a major shortcoming that made the algorithms unfit for implementation. The Complement algorithm seems most promising, but a full analysis has not yet been completed. The number of centers chosen is uncalculated, but in every other respect, the algorithm scores either well or exceptionally well.

High-Speed Multi-Component Flows

Student **Jeffrey W. Banks,** Rensselaer Polytechnic Institute

Mentor **William D. Henshaw,** CASC

We studied fluid flows with more than one component as described by the well-known multi-component Euler equations. It is a well-known phenomenon that numerical studies of these equations in conservation form can lead to unphysical pressure oscillations.

There are two main methods to combat this problem. The first is to investigate the system from a primitive variable formulation. This clearly necessitates a correction to bring the method back to a quasi-conservative form so as to maintain proper shock speeds. The other is to use some sort of conservative scheme on all but the fluid fraction variables and then advect those variables based on a special scheme that disallows the aforementioned pressure oscillations.

We present a third approach to this problem whereby one begins with a conservative method, but includes a source around contacts. This allows for all variables to be advanced through the same method while maintaining a conservative method in regions of the flow where there are no contacts between different fluids. We use an adaptive mesh refinement (AMR) scheme that is itself a constituent of an overset grid method and add the capability for moving grids. This allows our solver to operate on complicated geometries while maintaining efficiency through the AMR scheme.

We have demonstrated convergence through the method of analytic solutions and the so-called twilight-zone flows. We then performed calculations of the interaction of a shockwave with cylindrical bubbles of helium (*He*) and refrigerant-22 in order to show agreement with previous experiments and numerics. Then we used this solver to calculate the interaction of a shock with a special shaped piece of metal as in a shaped explosive charge.

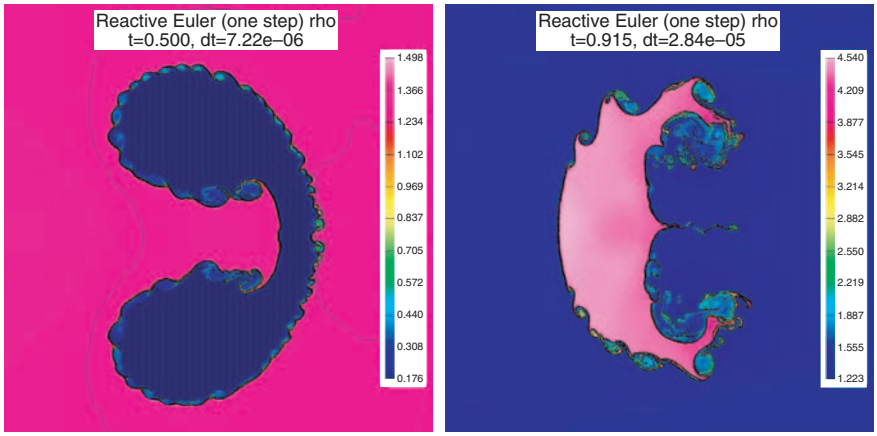

On the left, interaction of shock with He bubble, and on right interaction of shock with refrigerant-22 bubble. Both pictures have the shock running left to right, are at late time after the shock has passed, and clearly show the dominance of unstable features (Kelvin-Helmholtz, Rayleigh-Taylor).

TOC

Firewall Egress Filtering

Student **Lerone Banks,** University of California, Davis

Mentor **Jerry Rayome,** CSP

During my summer internship, I worked on several projects. My primary assignment was focused on presenting a solution for implementing egress filtering on the firewalls here at LLNL. Accomplishing this task required me to gather information about current network usage and requirements. From there, I explored firewall configuration details for the Green Network. With this information, I investigated current best practices in egress filtering and presented a solution for the preliminary steps of implementing egress filtering. Some of my other duties included evaluating various intrusion detection tools, making recommendations about their suitability for deployment and suggesting ways to secure the use of the domain name system (DNS).

Volumetric Mesh Parameterization Towards Slow-Growing Subdivision

Student

Janine Bennett, University of California, Davis

Mentor

Valerio Pascucci, CASC

Scientific simulations often generate unstructured mesh domains of arbitrary topology with data values sampled at the mesh nodes. However, a hierarchical, regular mesh structure is preferred for efficient analysis and rendering of the data. Subdivision methods are one of the most successfully used techniques in multi-resolution data representation for surface meshes. Initial research has been done to extend these techniques to the multi-resolution representation of volumetric data. For example, the Slow Growing Subdivision (SGS), first described in [1], is a subdivision method that generalizes to volumetric and higher-dimensional meshes. This approach currently works well for convex domains and regular curvilinear grids, where a mapping to a regular structure is provided explicitly.

This summer, we extended SGS so that it will work on arbitrary volumetric grids with non-convex boundaries. We use a region-growing technique that starts from single-cell sets, which are expanded as long as their structural properties are maintained in order to decompose the input domain into topological balls. We have developed a two-step process to map the resulting regions to an octahedral domain, i.e., we first map the boundary vertices and then the interior mesh vertices. To map the boundary, we determine edge paths that correspond to the edges of an octahedron. Using the method proposed in [2], we map the remaining boundary vertices to the faces of the octahedron. We are currently exploring methods to map the interior mesh vertices in order to obtain the necessary isomorphism from unstructured mesh connectivity to SGS connectivity.

Publications

[1] Valerio Pascucci, "Slow growing subdivision (SGS) in any dimension: Towards removing the curse of dimensionality," in *Proceedings of Eurographics* 2002, pp. 451–460, Saarbrucken, Germany, September 2002.

[2] M. S. Floater, "Parametrization and smooth approximation of surface triangulations," *Comp. Aided Geom. Design* 14 (1997), pp. 231-250.

Adaptive Algebraic Multigrid Preconditioners in Quantum Chromodynamics

Student **James Brannick,** University of Colorado at Boulder

Mentor **Rob Falgout,** CASC

Standard algebraic multigrid methods assume explicit knowledge of so-called *algebraically smooth* or *near-kernel* components, which, loosely speaking, are potentially large errors that correspond to relatively small residuals.

Typically, these methods automatically generate a sequence of coarse problems under the assumption that the near-kernel is locally constant. The difficulty in applying algebraic multigrid to lattice Quantum Chromodynamics (QCD) is that the near-kernel components can be far from constant, often exhibiting little or no apparent smoothness. In fact, the local character of these components appears to be random, depending on the randomness of the so-called "gauge" group. Hence, no *a priori* knowledge of the local character of the near-kernel is readily available.

Our work develops an adaptive algebraic multigrid (AMG) preconditioner suitable for the linear systems arising in lattice QCD. The method is a recently developed extension of smoothed aggregation, aSA, of Brezina, Falgout, MacLachlan, Manteuffel, McCormick, and Ruge, in which good convergence properties are achieved in situations where explicit knowledge of the near-kernel components may not be available. This extension is accomplished using the method itself to determine near-kernel components automatically by applying it carefully to the homogeneous matrix equation, $Ax=0$. The coarsening process is modified to use and improve the computed components. Preliminary results with model 2D QCD problems suggest that this approach may yield optimal multigrid-like performance that is uniform in matrix dimension and gauge-group randomness.

A Memory Insensitive Format for Out-of-Core Access to Unstructured Volumetric Meshes

Student **Steven Callahan,** University of Utah

Mentor **Valerio Pascucci,** CASC

Unstructured volumetric meshes are the preferred data format for scientific simulations. With the size and complexity of these meshes increasing every year, techniques have been developed to visualize the pertinent information quickly and efficiently. This visualization is a commonly performed preprocessing step in which a triangulated isosurface is extracted for a particular isovalue. For large meshes, this extraction is performed out of core. The work I have done introduces a global indexing scheme that accelerates the traversal of the data structure on disk by improving the spatial locality of the data. In addition, the data is converted to a coherent streaming format that facilitates IO-efficient, out-of-core algorithms.

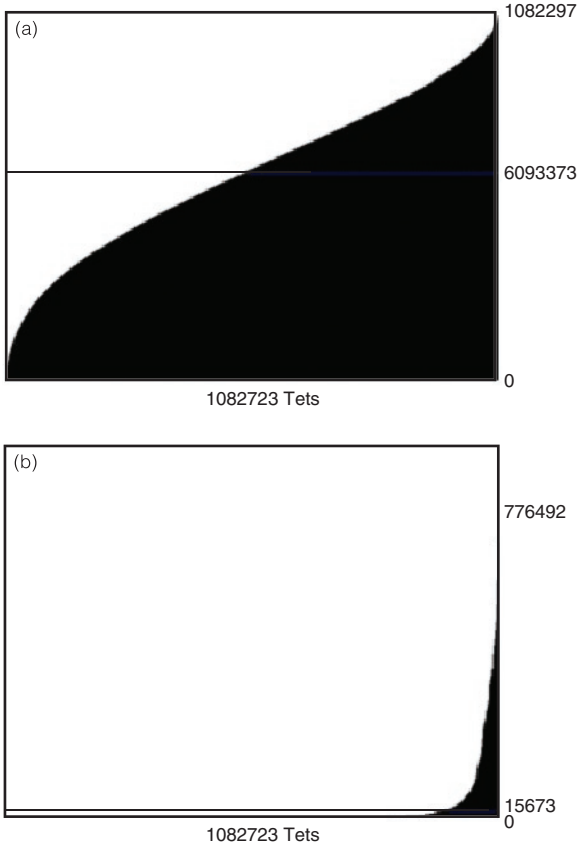

Figure 1. Comparison of original vertex order (a) and Lebesgue order (b)

The key idea is to order the vertices in such a way that minimizes the external memory accessing needed to traverse the volumetric cells. An out-of-core technique was developed to re-index the vertices using a Lebesgue space-filling curve and output the data in a streaming format. This new global indexing was then compared with previously proposed techniques using tools that displayed the IO cost of volumetric cell traversal. The contributions of my work include an out-of-core algorithm for re-indexing arbitrary unstructured volumetric meshes, an extension of streaming meshes for unstructured volumetric data, and a technique for analyzing the efficiency of the resulting mesh.

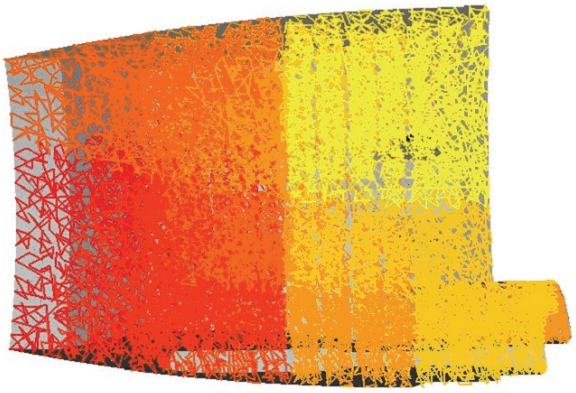

Figure 2. Lebesgue order imposed on the vertices of an unstructured mesh.

TOC

Discovery Center Display

Student **Jedidiah Chow,** Granada High School

Mentor **Jean Shuler,** CSG

The goal of our project was to produce an interactive program that visitors could use in order to inform them of the computing goals and accomplishments of the LLNL Computational Directorate.

To keep things simple and editable, only well-known programs were used to make the project. Microsoft Power Point is the basis for the display program, and it proved adequate for navigating through various topics, as well as being aesthetically pleasing. Basically, the program elaborates on supercomputing history at LLNL, visual display equipment, and the recent uses of the supercomputers. The project was completed under the supervision of Terry Girill and Bob Howe.

As with any program, periodical general maintenance and modifications will be necessary, and major updates will be accessed as time progresses and the information presented becomes out of date.

Incorporating Electrokinetic Effects into the EB Navier–Stokes Embedded Boundary Incompressible Fluid Solver

Student

Kevin Chu, Massachusetts Institute of Technology

Mentor

David Trebotich, CASC

Motivated by the recent interest in using electrokinetic effects within microfluidic devices [1], we have extended the EB Navier-Stokes embedded boundary incompressible fluid solver [2] to be able to handle electrokinetic effects. With this added functionality, the code will become more useful for understanding and designing microfluidic devices that take advantage of these effects (e.g., pumping and mixing).

Supporting the simulation of electrokinetic effects required three main extensions to the existing code:

(1) Addition of an electric field solver
(2) Development of a module for accurately computing the Smulochowski slip-velocity at fluid-solid boundaries
(3) Extension of the fluid solver to handle non-uniform, inhomogeneous Dirichlet boundary conditions.

The first and second extensions were needed to compute the electrokinetically generated slip-velocity at fluid-solid boundaries. The third extension made it possible for the fluid flow to be driven by a slip-velocity boundary condition rather than by a pressure difference between inflow and outflow.

We are in the process of verifying and validating the code. Figure 1 shows the time evolution of the horizontal component of the fluid flow around a single cylindrical post where the flow is driven by electrokinetic slip at the surface of the posts and the top and bottom walls. In the future, we hope to use the code to simulate electrokinetically driven flows in novel microfluidic devices being explored for biological applications.

References

(1) T. M. Squires and M. Z. Bazant, "Induced-charge electro-osmosis," *J. Fluid Mech.*, 509 (2004), pp.~217–252.

(2) D. Trebotich, "Working Notes for Higher-Order Projection in Embedded Boundary Framework," private communication.

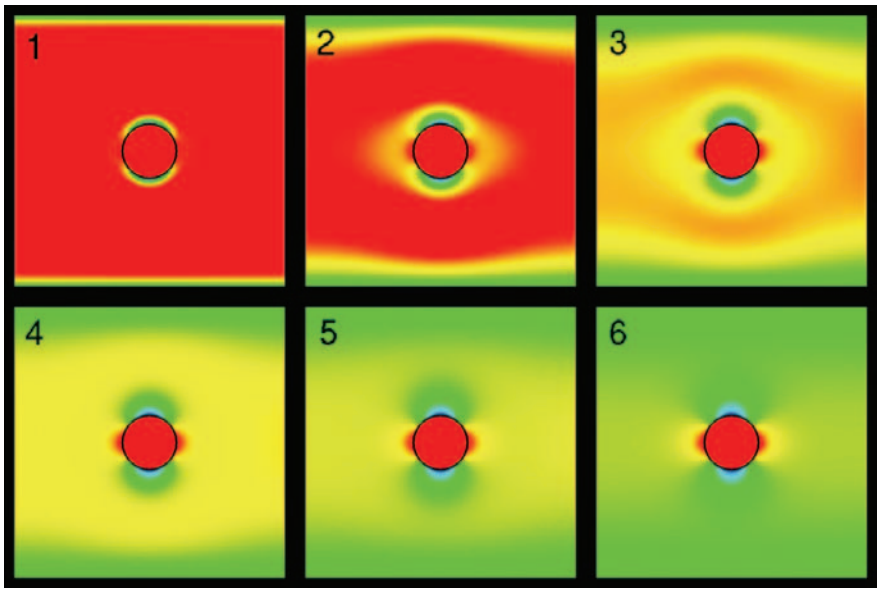

A flow driven by electrokinetic slip moves around a cylindrical post

Applying Morse Theory to Computational Data Sets

Student

Kree Cole-McLaughlin, University of Utah

Mentors

Valerio Pascucci, CASC **& Terence Critchlow,** CASC

Morse theory is a powerful tool that connects the theory of real-valued functions over a manifold to the topology of the manifold. Recently, computer scientists have begun to develop computational tools for generating topological characterizations of scientific data using Morse theory. The contour tree, for example, characterizes how the connected components of the level sets evolve as the function value increases.

Another important characterization is the Morse-Smale complex, which decomposes the data set into regions of uniform gradient flow. This is a more complete characterization than the contour tree and is in some sense the essential application of Morse theory. Algorithms for computing the Morse-Smale complex over the 2- and 3-spheres have been developed within the last few years.

The major problem in applying Morse theory to computational data sets is that the theory was developed for smooth manifolds and differentiable functions. Previous techniques simulate the differentiability of the function and use combinatorial data structures to keep track of several types of degeneracies that can arise. In two dimensions this does not cause much of a problem, however, for three-dimensional data sets, it adds a great deal of complexity to the description and implementation of the algorithm.

Recently Prof. Robin Forman developed a discrete version of Morse theory that can be thought of as generalizing the smooth theory. Forman's version of the theory has faced some problems in finding applications, due mainly to the fact that functions are defined in an apparently unnatural way. However, over the summer, we developed a technique for applying Forman's theory in such a way that we can prove that the results are equivalent to the previous techniques.

The benefits of our approach are many. Most significantly is the simplicity of the algorithms, which require no additional data structures for dealing with degeneracies. In some sense, we push all the complexities of the previous approaches into the theoretical framework and deal with them in the proofs of theorem. Our algorithms, on the other hand, remain simple. Additionally, our approach has the potential for being extended beyond three-dimensional data sets, where Morse theory can play a vital role in understanding the data, since visualization is nearly impossible. Finally, we do not need as many restrictions on the underlying manifold as the previous techniques, which only work on spheres. This technique is still being worked out, but our work this summer has proven that it is valid.

TOC

Geometric Multigrid for Variable Coefficient Maxwell's Equations

Student **Dylan Copeland,** Texas A&M University

Mentor **Barry Lee,** CASC

We are interested in numerically solving Maxwell's equations of electromagnetics with variable coefficients. Variable coefficients are difficult because the bilinear form that corresponds to the variational formulation has a large near-nullspace consisting of more than just gradients. Consequently, methods such as Hiptmair's hybrid smoother are ineffective on some components of the error in the near nullspace.

We implemented a multigrid method that uses interpolation operators designed so that near-nullspace components are locally interpolated from the coarse grids. To approximate curl-free components of the solution, we solved the corresponding scalar Poisson equation with the Maxwell equation simultaneously. Interpolation operators were then needed for scalar and vector finite element spaces. For the scalar spaces, we used BoxMG interpolation, which is known to be very efficient for Poisson equations on structured grids. The interpolation for the Nedelec vector finite element space was constructed locally so that coarse grid functions were interpolated to fine grid functions in the near-nullspace of the operator associated with the variational formulation. This was accomplished by the inversion of local element matrices, which is quite inexpensive.

To test the efficiency of our multigrid method, we performed numerical experiments with coefficients containing oscillations and jumps. The number of iterations required by the V(1,1) cycle to reduce the L2 norm of the residual by eight orders of magnitude was approximately 40% less than that of Hiptmair's multigrid method. Thus, a significant improvement in the convergence rate is achieved with our method.

TOC

Web Designer, Program for Climate Model Diagnosis and Intercomparison

Student **Steven Davis,** California State University, Hayward

Mentor **Charles Doutriaux,** EEBI Division

My assignment was to make the previous version of the PCMDI Web site http://www-pcmdi.llnl.gov/old easier to use and more aesthetically pleasing to the user. This was accomplished by creating a consistent design for the entire site, recreating the directory structure to make it more understandable, and implementing new features that allow users to navigate more easily, such as a drop-down menu system, breadcrumb navigation, and a search box on every page.

These features, and many smaller details, were added and changed to convert the PCMDI Web site to what is today http://www-pcmdi.llnl.gov. I will continue to work one day a week during the school year to add more functionality to the site, such as database-driven dynamic pages, as well as the ability for group members to update certain pages without compromising the design.

Edge Betweenness Properties in Complex Networks

Student

Stanko Dimitrov, University of Arizona

Mentor

Edmond Chow, CASC

Vertex betweenness, the number of shortest paths that pass through a vertex, has been studied extensively in social networks; the notion of vertex betweenness was later extended to include edge betweenness. These measures can be thought of as ranking the vertices (in the case of vertex centrality) and the edges (in the case of edge centrality) in order of importance. This importance can help people to prevent epidemics from spreading by removing people with the highest rank, or seeing that one edge is a bridge between communities (a cluster of vertices). However, no formal, mathematically formulated definition and exact algorithm for edge betweenness was presented.

We initially formally defined edge betweenness and presented a fast algorithm for computing edge betweenness by modifying the current fast algorithm for computing vertex betweenness. With our formal definition of edge betweenness, we investigated how edge betweenness behaves in different types of complex networks, where a complex network is a large graph with a specific probability distribution on its vertices that is used to determine the edges of a graph. The vertex probabilities help determine which two vertices make up an edge, meaning the vertex with the highest probability is more likely to make up an edge than a vertex with the lowest probability.

In our experiments, we had three types of complex networks—random, spatial, and scale free. We also had one real network that was part of the network of routers that make up the Internet. We observed how the measure behaves in our models and how it behaves in the real network. We need to know how our measure behaves so we can classify our models. With this classification, we hope to be able to say something about the structure of a real network given that we know how our measure behaves in that network.

One drawback to edge centrality is it takes a long time to compute, which means we are interested in finding other measures that can approximate edge centrality. With these approximations, we can quickly rank the edges and begin making statements about the structure of our complex networks.

TOC

Element Agglomeration AMGe Solvers for Unstructured Finite Element Problems

Student **Veselin Dobrev,** Texas A&M University

Mentor **Panayot Vassilevski,** CASC

Unlike the standard (geometric) multigrid algorithms, which require a sequence of coarse operators and interpolation/restriction operators, the AMG methods only require a fine-grid matrix as input, which is used to construct the coarse level matrices. Similarly, the element agglomeration AMG (AMGe) methods, which can be used when the finite-element method (FEM) is used for discretization, also use fine-grid information only—the element matrices combined with topological information for the FE mesh and discretization space. Therefore, these methods can be used for problems on unstructured meshes and complex domains where standard multigrid is not applicable.

Based on the Texas A&M finite element package AggieFEM, a code was developed for discretizing various types of PDEs in 2D and 3D, including second-order elliptic equations, linear elasticity, etc. The code uses arbitrary triangular (2D), tetrahedral and hexahedral (3D) meshes, allows for global (and in some cases local) refinement, and supports several types of finite elements—linear and quadratic conforming, and linear nonconforming. The generated discretizations were used as input to an AMGe library written by Dr. Vassilevski for testing and experimental purposes and to evaluate the solvers' performance. The developed code also provides visualization for the obtained solutions and element agglomerations (generated by the AMGe code). As a result, improvements and enhancements were made to the library. In particular, a new agglomeration algorithm was developed based on the graph partitioning library METIS.

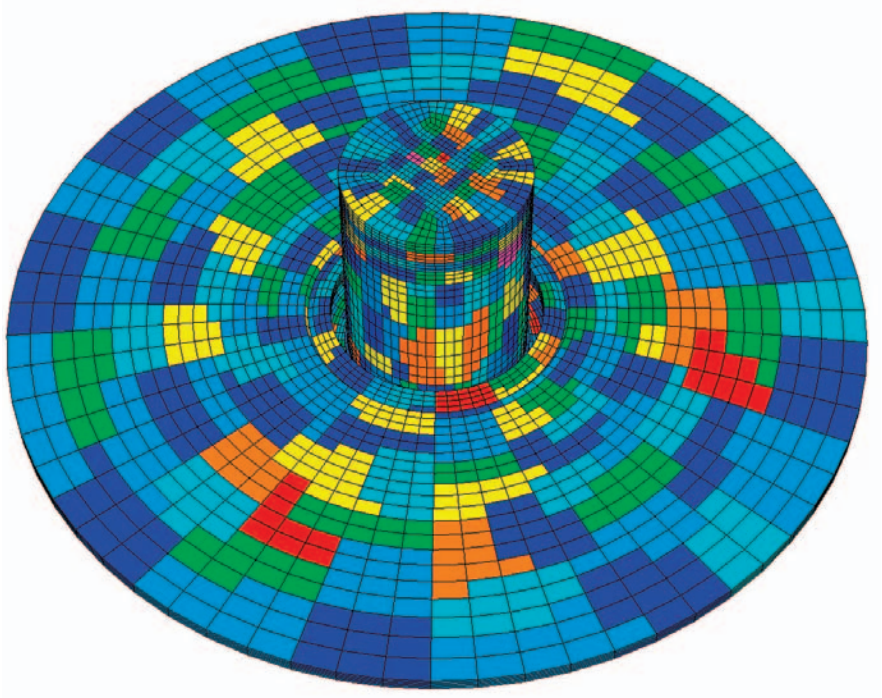

Element agglomeration in 3D

TOC

Parallel Implementation of a FETI–DP Method for Elliptic Problems with Mortar Finite Element Discretization

Student

Nina Dokeva, University of Southern California

Mentor

Panayot Vassilevski, CASC

This project focused on elliptic problems with discontinuous coefficients discretized by Finite Element Method (FEM) on non-matching triangulations across the interface using the mortar technique. The resulting problem is solved by a Finite Element Tearing and Interconnecting Dual Primal (FETI-DP) method proposed by M. Dryja and O. Widlund, and generalized by adding multiple scalings in the preconditioner to problems with discontinuous coefficients by N. Dokeva, M. Dryja and W. Proskurowski. The convergence has polylogarithmic dependence of the mesh-sizes. Moreover, it is independent of the number of subdomains which makes the algorithm suitable for parallel implementation with one processor working on each subdomain. The parallel code was developed using C and MPI. It ran successfully on 256 processors.

After domain decomposition, the problem is divided into three types of tasks:

(1) Solvers on the subdomains (with different meshes of discretization) which run individually and in parallel

(2) A problem on the interfaces between the subdomains which can be solved in parallel with only a few global communications

(3) A problem on the cross-points between the subdomains, which is a global task

Scalability when increasing the number of processors was only acquired after optimizing the individual parts of the algorithm, paying the most attention to the problem with the Schur complement matrix on the cross-points between the subdomains. Ultimately, the algorithm proved to be scalable when increasing the number of processors, with only a minimal overhead in communications, which significantly decreased the time for executing the program compared to the sequential algorithm.

Future work will aim at improving the cost increase when refining the mesh, making further improvements in the preconditioner, and considering different mortar functions.

TOC

Graph Viewer Improvements

Student **Emily Eder,** UCLA

Mentor **Robert Fernandes,** IOAC

My assignment was to make improvements to the Java-based graph viewer currently in use at the Lab. The first part of my assignment was to make a URL Mapping Panel. This was accomplished by creating a Java dialog box to input corresponding Web locations and file locations and addressing the need to store and retrieve corresponding Web locations and file locations.

The second part of my assignment was to conduct research into graph layout algorithms for the graph viewer. My research focused on force-directed algorithms and Java implementations of these algorithms, including the Geometric Efficient Matching (GEM) algorithm, the spring embedder algorithm, the simulated annealing algorithm, and graph drawing by stress majorization. The goal of the project was to identify the most efficient force-directed algorithm, which would then be implemented in the graph viewer within the next six months. The GEM algorithm was identified as one of the most effective force-directed algorithm.

Platform Independence for CALE: wxWidgets and GNU Autotools

Student

Christopher Egner, Rochester Institute of Technology

Mentor

Paul Amala, A Division

My work centered around two objectives, implementing a graphical user interface for C-language Arbitrary Lagrange Eulerian (CALE) using wxWidgets and reimplementing CALE'S configuration and build system using GNU Autotools. CALE'S user interface was based in X11 using a single-threaded architecture. As such, there were no concurrency issues, and the X11 API calls, while wrapped by other functions, were more coupled to the physics code than would be standard in a modern Model-View-Controller design.

In order to combine the wxWidgets and CALE architectures, a multi-threaded model was chosen last summer. This introduced two related problems. First, CALE'S code was not designed to be called in more than one thread, so it is not well-suited to a normal multi-thread model. Second, wxWidgets is not thread-safe, imposes the requirement that all graphics calls (an ill-defined term) be made from the main thread, and assumes that the program's execution architecture can be governed by its event-processing loop. It is these architectural issues that are at the foundation of many of the problems I encountered, and some remain unresolved. Completed work includes an implementation of buffered graphics, improved efficiency, increased stability, patches to wxWidgets, and a change in template models for wider compiler compatibility.

CALE'S former build system was based on an M4 script that produced a single makefile to build either the high-explosive or inertial-confinement fusion version of the code with manual platform detection and manual dependency tracking. Reimplementing the configuration and build system with GNU Autotools resulted in an automated and more customizable configuration system, as well as a more maintainable, faster build system with automated dependency tracking. Combined, the new systems configure and build not only the high-explosive and inertial-confinement fusion versions of CALE, but also the project's supporting tools and documentation.

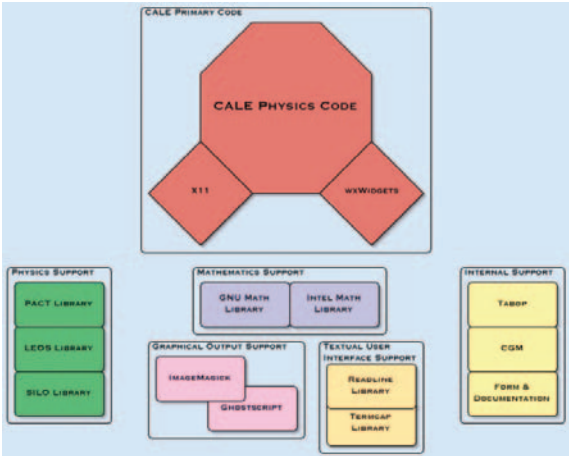

Figure 1. The structure of CALE and its supporting libraries and tools, grouped by purpose. CALE can be built with either the X11 or wxWidgets GUI and with any supporting libraries available.

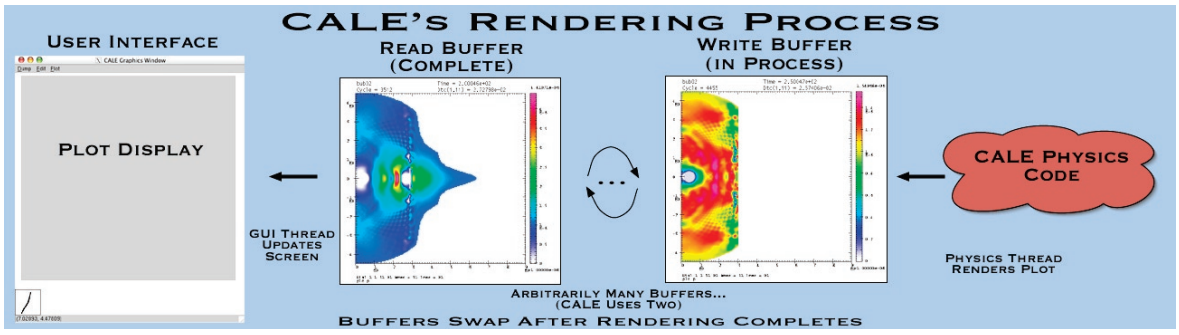

Figure2. CALE'S buffered graphics rendering system. CALE can be used to render animations as it computes. However, animation can lead to flickering on the screen as the program redraws. Thread-safe, rotating buffers assure that CALE can always refresh the screen, even as it works on the next plot. By storing the last rendering, CALE can load it quickly to the screen in order to repair damaged areas of the screen without having to recompute. By computing the entire plot first, then drawing a bitmap rapidly to the screen, flicker is appreciably reduced and drawing is significantly faster.

TOC

Wavelet-Based Opacity Data Compression

Student **Christopher Elofson,** University of Arizona

Mentor **Vijay Sonnad,** DCOM

Astrophysical opacity data sets are large and often contain many spectral lines. A compression algorithm capable of accurately representing both the peaks and depths between them (related to a quantity known as the Rosseland mean) is desirable, as this will provide greater computational efficiency.

While compression could be implemented using a Fourier transform, the Fourier method cannot be inexpensively adapted for data with many peaks, as the total number of samples must be increased to improve local resolution.

A better approach to opacity data compression is to use Discrete Wavelet Transforms, which differ from Fourier transforms in that they can use any function as a basis function—Fourier can use only sinusoidal functions. Wavelets are not only a better match to data's end behavior, they are also a much less expensive, yet equally accurate, representation of spectral lines, made possible by wavelets' capability of local resolution adjustment.

Through the testing of six varying opacity data sets, the Discrete Wavelet Transform proved to be a very effective way to increase manageability of opacity data. In five of six cases, the linear transform provided the best results, with compression rates varying from 45–99% (three of six cases yielding rates over 98%) while conserving the Rosseland mean within 0.00001%—the Haar basis varied between 38–96% compression. In the case where many spectral lines were densely packed over the entire domain, the Haar basis performed better in data representation, producing a lower rate of error in the Rosseland mean as well as lower absolute error in comparison to the linear compression.

TOC

ODE Visualization and Archive Tool in Python

Student

Dayo Esho, University of California, Berkeley

Mentor

Steven Lee, CASC

The Runge-Kutta method is a numerical technique used to approximate the solution of ordinary differential equation (ODE) initial-value problems. We developed a visualization and archiving tool to interface between the data output of advanced ODE solvers. In particular, we studied higher-order, adaptive-stepsize, explicit Runge-Kutta methods by Dormand and Prince (D&P) and Calvo, Montijano, and Randez (CMR). By analyzing the visual output from these solvers through our tool, we can compare solutions between solvers, determine their relative accuracy and efficiency, and monitor solver performance statistics. Archiving these results in a Web interface allows for rapid regeneration of particular experiments and easy modification of parameters to extend an analysis of a particular solver-problem combination. Our work focused on the ability of these solvers to handle stiff ODEs.

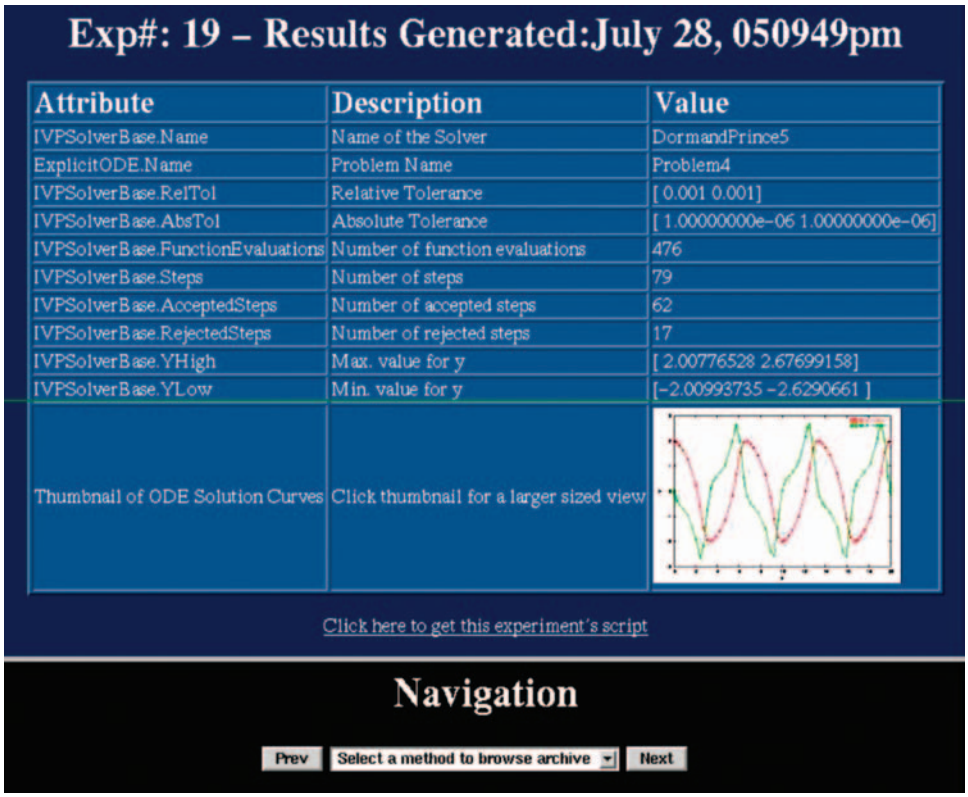

Screenshot of the ODE Viz/Archive Tool

2D Numerical Modeling of Soap Film on Overture

Student **Pak-Wing Fok,** Massachusetts Institute of Technology

Mentor **Petri Fast,** CASC

My assignment was to learn how to use Overture, a set of C++ Libraries that facilitate the solution of partial differential equations in complex geometry, and then use that knowledge to code up a solver for a set of soap film equations recently derived by Dr. Petri Fast. Following this, we simulated the short-time evolution of a soap film, in the middle of which a solid cylinder was immersed (see below). I also performed some error analysis and convergence studies to validate my numerical solution.

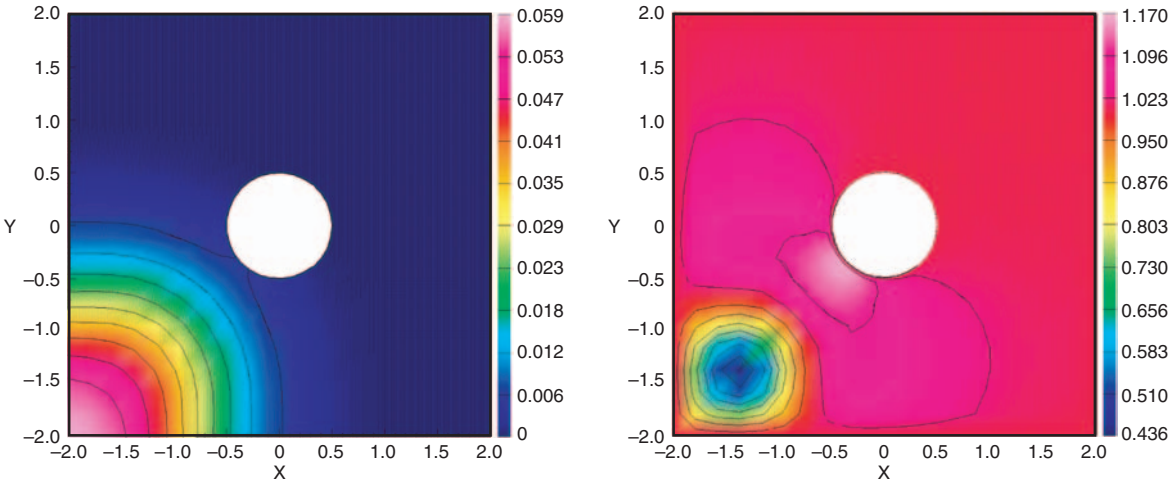

Two snapshots from my numerical runs. The one on the left shows surfactant concentration; the other shows the thickness of the film.

TOC

Implementing Cray Pointers in the GNU gfortran Compiler

Student

Camille Fournier, University of Wisconsin

Mentor

Mike Kumbera, DCOM

During my assignment, I began the implementation of Cray pointer support in the GNU gfortran compiler. Specifically, I modified the parser to recognize Cray pointer statements, and modified the gfortran frontend to produce code that correctly represented these statements to the gcc backend so that it could compile them. This process included creating a new data type for Cray pointer data, adding command-line arguments to specify whether these pointers were used in the program, and adding intrinsic functions that were required in order to use pointers of this type.

Performance Analysis of Monitoring and Information Systems Using NetLogger and an Investigation into P2P Overlay Topologies

Student

Jeffrey Freschl, University of Wisconsin, Madison

Mentor

John Johnson, DCOM

In collaboration with members of Globus at Argonne National Laboratory (ANL), we analyzed and compared the performance of three monitoring and information services—the Relational Grid Monitoring Architecture used in the European Grid test bed, MDS2 used in the Globus toolkit, and Hawkeye used in Condor pools. Our analysis involved instrumenting the source code with NetLogger calls to locate performance bottlenecks. With the instrumented tools, we simulated up to 600 clients, 600 monitored machines, and 80 sensors. Our results indicate that running the main components of such systems at well-connected sites (i.e., the network is a bottleneck) and caching the monitoring information (i.e., activating an information sensor is a bottleneck) can provide significant performance improvements. In addition to completing a poster presentation, we are preparing a journal paper.

My work on peer-to-peer overlays involved a survey of the current work on the overlay mismatch problem. The main methods include using landmarks, clustering to partition the overlay, and locally optimizing a peer's connections to its neighbors. The main use of current P2P networks involves file sharing (e.g. Kazaa), and the overlay mismatch does not pose much of a performance problem. However, someday we may be able to utilize the Internet for far more compute-intensive, parallel algorithms (not necessarily restricted to embarrassingly parallel algorithms) where communication can be a significant bottleneck. Would it be possible to use a P2P network for tightly coupled simulations? Perhaps not to that degree, but high-throughput computing could utilize a P2P network as a computational resource.

Bringing Order to CHAOS

Student **Evan Geller,** Summit Preparatory High School

Mentor **Albert Chu,** HPSD

My major project focused on Cerebro, a host monitoring system. I coded the networking piece and config file parsing that enabled system administrators to work without having to make a config file for each monitored node.

I also created reproducers and integrated them into our test suite to chronicle the following issues:

- A bug found in the cpickle function in Python which caused a memory leak
- A bug found in the Network File System (NFS) that resulted in chmods inconsistency across nodes
- A bug, found in the getrusage system call that allowed system time to be reported incorrectly
- A bug that was found in the interaction between rsync and NFS that caused stale file handles.

In addition to this work, my mentor found a test to assess various aspects of NFS, which I was able to integrate into our test suite and report on NFS reactions. Our group needed a script that would monitor how NFS reacted to various commands. To address this need, I wrote a script that wrapped around the given command and observed how NFS reacted.

Assessing Performance of Hybrid MPI/OpenMP Programs on SMP Clusters

Student

Tobias Gradl, Technical University of Munich

Mentor

Edmond Chow, CASC

The SMP cluster architecture is the prevalent type of computer system in today's high performance computing landscape. Two strategies are popular for programming these computers. In the *pure MPI approach,* message passing is used for all data exchange, even between processors within a node. In the *hybrid approach,* messages are only passed between processors on different nodes, while processors within a node share data through multithreading.

While hybrid programs tend to demand more coding effort, they promise to be faster than pure MPI programs. Practical experiments show, however, that in many cases both strategies are equally fast; some even show that pure MPI is faster. Some of the reasons for this behavior are well understood, while others are not. The insights gained from the comparison will help explain the behavior of large real-world applications. They can also guide programmers in the design of new applications.

We used a simple but realistic benchmark, the matrix-vector product, to compare the two programming strategies. The simplicity of the matrix-vector product helped us analyze and understand its behavior in detail.

We ran the benchmark program using different problem sizes and numbers of processors and took timings of all code sections. The characteristics of the collected timings were then explained by various environmental factors like cache size, memory bandwidth, network bandwidth, etc. Knowing how these factors affect the program performance enables us to extrapolate relevancy to other applications and system configurations.

TOC

Censorship Resistance

Student **Rachel Greenstadt,** Harvard University

Mentor **Terry Brugger,** NAIC

I participated in three main projects during the Lab's summer program: Circumventing Government Firewalls, LOCKSS (Lots of Copies Keep Stuff Safe), and Indexing Hacker Tools.

Many countries around the world, such as China, Singapore, Iran, and Saudi Arabia, have censored access to the Internet. Traditionally, these firewalls have been circumvented with proxy servers, which are also eventually censored. This project explored efforts to win this arms race and identified key bottlenecks where more work is needed. We concluded that circumvention of such systems is easy to accomplish on a small scale, but difficult on a large scale. The hardest problem seems to be publicizing information about the circumvention system without having the system blocked.

LOCKSS is a peer-to-peer system that allows digital libraries to collaborate and archive their data. It aims to protect information stored on a single system from being distorted due to bit rot, targeted censorship, or physical destruction. The LOCKSS project includes participants at multiple universities and a deployed beta system with 80 participating libraries. Libraries periodically poll a sample of peers in the system to determine if their copy of the document is good. Pollees consist of a combination of trusted "friends" and other discovered peers. In the event there is no clear supermajority of peers, an alarm is called as an intrusion detection measure. We conducted simulations of a network of 1,000 peers wherein each peer conducted 3,600 polls to determine the most effective, practical mechanism for a peer to respond to alarms. This mechanism proved to be a procedure in which peers who experienced an alarm contacted and healed their friends.

For the third project, I anonymously surfed a variety of hacker and security Web sites in order to create an HTML table describing the range of hacker tools currently available in the open source community and indexed them by category.

TOC

Hierarchical Morse–Smale Complexes

Student **Attila Gyulassy,** University of California, Davis

Mentor **Valerio Pascucci,** CASC

A Morse–Smale complex is a structure that represents the topology of a data set. Each cell of the complex consists of a minimum, a set of 1-saddles and 2-saddles, and a maximum. Connectivity of these critical points is also stored in the cell. The Morse–Smale complex can be simplified for use in applications such as topological smoothing or multi-resolution viewing.

We developed an algorithm for simplifying Morse–Smale complexes that relies solely on combinatorial decisions, thereby avoiding numerical instability. Simplification of a Morse–Smale complex is possible through the cancellation of critical point pairs. There are two types of cancellations—saddle–extrema cancellations and 1-saddle–2-saddle cancellations. We developed rules for determining when a cancellation is valid in a complex, and for reconnecting the simplified complex.

A multi-resolution representation of a Morse–Smale complex can be achieved through a simplification hierarchy. We presented a set of rules that govern valid simplifications of a complex and developed a hierarchy based on the independence of those cancellations. We declare two cancellations to be independent when their affected areas in the Morse–Smale complex do not overlap. We represent the hierarchy as a directed acyclic graph that encodes the independence of cancellations. A multi-resolution reconstruction of the complex can be attained by cutting the hierarchy graph.

Future work involves extracting a Morse–Smale complex from a scalar data set and applying the simplification hierarchy to topology-based smoothing of actual data sets.

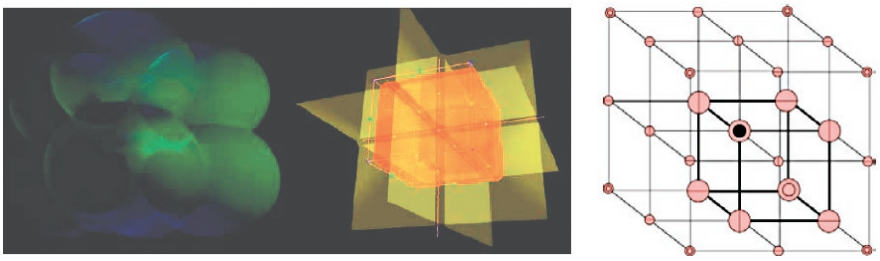

Figure 1. An isosurface of a simple data set (left), the corresponding stable and unstable manifolds, and the associated Morse-Smale complex.

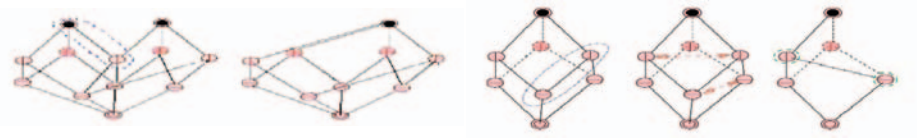

Figure 2. A saddle-maximum cancellation (left) and a 1-saddle–2-saddle cancellation.

MPI Profiling

Student **Daniel Han,** University of Southern California

Mentor **Terry Jones,** Services & Development Division

Message Passing Interface (MPI) is the de facto standard in writing massively parallel programs and has been embraced by the scientific community for use in solving a variety of problems. Performance is often a crucial factor, especially of grand challenge problems. Although there have been many studies on the scalability of these applications, there have not been many that focused on the specific types of MPI calls being made and their relative importance. Using a profiling tool called mpiP, I surveyed a large spectrum of parallel scientific applications and analyzed the results.

Parallel Analysis of Asymmetries in Symmetric Simulations in ALE3D

Student **Kevin Hoffman,** Brigham Young University

Mentor **Robert Cooper,** DCOM

The unstructured finite-element mesh hydrodynamics simulation code ALE3D efficiently scales to problems with millions of zones. Validation of the correctness of the code as it is developed is critical to maintaining the accuracy of the simulation. One way to validate the code is to simulate a symmetrical problem and ensure that a reasonable level of symmetry is maintained in the solution fields.

I designed and developed a parallel analysis tool to determine the amount of asymmetry at each zone and node in the mesh. Some challenges included efficiently determining communication needs between processors, efficiently pairing each point in the mesh with its symmetrical counterpart, and dealing with vector-based variables.

I further enhanced the algorithm to support taking the difference between two fundamentally different meshes, such as problems meshed at different resolutions or a mesh partitioned using different methods. This allowed me to verify the preservation of the order of operations using different domain partitionings.

I also developed a parallel statistics calculation framework that calculates various statistics over different subsets of the mesh—region, domain, problem, etc. The framework allows for easy addition of new statistics.

Finally, I researched how to redesign the domain partitioning code to meet the future needs of the ALE3D team. I created a flexible design that abstracts the interface between the mesh and the partitioner, allowing for various mesh sources, chaining of partitioners, refinement of partitioning, and dynamic repartitioning. I also wrote a partial implementation showing the validity of the design.

Combinatorial Feature Extraction for a Streaming Framework

Student

Taylor Holliday, University of California, Davis

Mentor

Valerio Pascucci, CASC

We devised and implemented combinatorial algorithms for constructing discrete Morse functions and Morse complexes on 2D and 3D simplicial complexes. Based on spanning trees, this combinatorial approach is easy to implement, amenable to a streaming multiresolution representation of the data, avoids the numerical instability of previous approaches, and scales to large data sets.

Morse complexes are useful for multiresolution visualization, topological smoothing, and mesh reparameterization. They simplify a data set while preserving extrema.

Our approach builds a discrete Morse function from a piecewise-linear function on the vertices. The algorithm partitions the data set into subcomplexes, which are ordered by function value, and builds minimal Morse functions on each subcomplex. Together, these minimal functions form the Morse function for the entire complex.

Future work includes adapting these algorithms to an out-of-core hierarchical representation of the data set, such as the Z-order space-filling curve.

The figure illustrates the 2D algorithm run on a trivial simplicial complex. A minimum (yellow), a maximum (magenta), and six saddles (cyan) are depicted as colored squares. The magenta and yellow lines indicate paths of steepest ascent and decent (respectively) from the saddles. These lines are the 1-cells (edges) of the Morse complex.

TOC

Virtual PCR (vPCR)

Student

Gary Hon, University of California, San Diego

Mentor

David Clague, EETD

I worked with University of Virginia summer student John Frankovich to integrate the nonlinear equation solver KINSOL into vPCR. The overall goal was for vPCR to solve a system of free energy and conservation equations that were both highly coupled and nonlinear. This involved deriving and experimenting with half a dozen different approaches. The final approach involved improving KINSOL to support arbitrary precision arithmetic using Lawrence Berkeley National Laboratory's ARPREC package and incorporating a free energy error term.

I improved vPCR by making simulations of higher-order reactions possible for both kinetics and thermodynamics codes. The original code only supported bimolecular reactions where each reactant was a monomer. The extended code does not have this limitation. This improvement involved completely restructuring the underlying data structures and the higher-order codes that use them.

Furthermore, I added the ability to simulate polymerase extension of annealed monomers. This purely kinetic simulation involved finding extension products and then running Gillespie's stochastic chemical kinetics to form products.

Lastly, I tested and verified the functionality of all of the aforementioned modifications.

A Two-Grid Method for Radiation Diffusion

Student **Jason Howell,** Clemson University

Mentor **Carol Woodward,** CASC

Applications of solvers to nonlinear radiation diffusion problems require repeated solutions of a linearized version of the problem, and these linear problems are often the bottleneck for the entire solution process. Our approach is a two-grid scheme in which the problem is solved very accurately on a coarse mesh followed by a single solver of a linearized version on the original problem mesh. The motivation for the method is that the coarse-mesh solution will capture the nonlinear behavior of the problem, and the fine grid solver will refine the solution on a precise scale, resulting in a faster overall solution process. This method was successfully applied to a similar problem arising in ground water applications by Wu and Allen in 1999. The main theoretical development for this method comes from Xu (*SIAM Journal of Numerical Analysis*, 1996), and the extensions to our nonlinear system come from Dawson, Wheeler, and Woodward (*SIAM Journal of Numerical Analysis*, 1998).

This method has been implemented at Lawrence Livermore National Laboratory using PVODE as the ODE time integrator and transport3d as the radiation diffusion model driver. Investigations of implementation issues addressed this summer included the transfer of solution data between the coarse and fine grids, and proper application of boundary conditions to the corrections solved for each grid. Many numerical experiments were run on linear and nonlinear problems, and the performance of the two-grid solver was compared to that of a standard nonlinear solver. Different choices of physical parameters provided a range of nonlinearity for our tests.

TOC

DHS Countermeasures Test Bed – Command Center

Student **David Hoyt,** Brigham Young University

Mentor **James Schek,** NAIC Division

The initial goal of this summer project was to design an online interface to the advanced radiation detector's embedded HTTP server. This included working on the layout, design, image editing, JavaScript programming, and flow. The design was presented to the end users, who made suggestions for minor changes, and it was unanimously accepted.

Thereafter, threading and Win32 API errors necessitated alterations to the existing serial port interface code. I was able to successfully implement a testing program and provide a new solution to the problem by rewriting major portions of the existing classes. I conducted research and testing to understand important programming concepts, such as threading and asynchronous input and output.

I then began to implement a security model in Java for the Command Center utilizing Java Data Objects (JDO) and a Pointbase database. As part of my work, I had to become familiar with JDO, including its implementation, metadata files, and integration with relational databases. I created a custom tag library for use in Java Server Pages (JSP) to complement the security model. It was then necessary to unify and centralize the logins.

PerfTrack

Student

Kevin Huck, University of Oregon

Mentor

Brian Miller, CASC

Our goal was to create a tool to help scientific programmers answer difficult questions about application performance, given that the source code, build parameters, runtime environment, and hardware vary over time. PerfTrack was developed to explore technologies in parallel performance measurement, modeling, analysis, and prediction. Performance data and the associated environment data is being stored in a relational database. This database provides a foundation to build analysis tools that are scalable to large numbers of threads (over 1,024) and that are capable of comparing multiple executions. The tools will be automated to gather, store, and analyze data, in order to encourage their use in the software development cycle.

In addition, I researched the use of data mining (cluster analysis) techniques on performance data by applying techniques used in phase analysis of applications to parallel performance data. As a result of this research, some initial comparisons between threads in a single execution were made. The results of the analysis were similar to other studies using cluster/factor analysis on the same application.

A Multilevel Preconditioner for a PDE-Constraint Optimization Problem

Student **Lukas Jager,** Universität Bonn, Germany

Mentor **Radu Serban,** CASC

In this project, we considered the problem of determining a source in a time-dependent, advective-diffusive transport equation by given measurements of the concentration at certain locations. This problem is formulated as a constraint optimization problem where the objective function represents the difference between these measurements and the model prediction. The so-called reduced-space approach allows the reformulation as an unconstrained least squares minimization problem, which is typically solved by a Newton-type method. This entails the solution of linear systems involving the Hessian and the gradient of the objective function.

Adjoint sensitivity allows the computation of the gradient and the action of the Hessian on some vector with the same cost as the solution that the model system requires. But since the Hessian is not assembled, an iterative method for the solution of the linear system has to be applied. Due to the large condition number of the matrix, classical iterative methods provide only slow convergence. The use of a preconditioner can improve the convergence by changing the properties of the system matrix.

We implemented a multilevel preconditioner that consists of one or more coarser discretizations of the source. To solve the preconditioning system, it applies a full multigrid V-cycle. On the coarsest level, the discretization allows the assembly of the Hessian and solves the system by inverting the matrix. Our multilevel preconditioner enhances the convergence of the iterative solver and hence reduces the total number of iterations during the minimization.

The project will be continued by improving the performance of the preconditioner and considering different types of source discretizations that allow, for example, the determination of a time-dependent source.

Integration of Multivectors to Hypre

Student
Abram Jujunashvili, University of Colorado at Denver

Mentor
Charles Tong, CASC

The main project focused on the integration of multivectors to Hypre. To accomplish this, we created new data structure—parallel multivectors and the functions that are necessary to use them in Hypre. First, we built the sequential realization of several functions:

hypre_SeqMulti_VectorCreate
hypre_SeqMulti_VectorDestroy
hypre_SeqMulti_VectorInitialize,
hypre_SeqMulti_VectorSetDataOwner
hypre_SeqMulti_VectorSetConstantValues,
hypre_SeqMulti_VectorSetRandomValues
hypre_SeqMulti_VectorCopy
hypre_SeqMulti_VectorScale
hypre_SeqMulti_VectorAxpy
hypre_SeqMulti_VectorInnerProdGram
hypre_SeqMulti_VectorMultiScale
hypre_SeqMulti_VectorMultiAxpy
hypre_SeqMulti_VectorInnerProdDiag.

Based on these functions, we then constructed the parallel functions for parallel multivectors:

hypre_MultiParVectorCreate
hypre_MultiParVectorDestroy
hypre_MultiParVectorInitialize
hypre_MultiParVectorSetDataOwner
hypre_MultiParVectorSetPartitioningOwner
hypre_MultiParVectorSetMask
hypre_MultiParVectorSetConstantValues
hypre_MultiParVectorSetRandomValues
hypre_MultiParVectorCopy
hypre_MultiParVectorScale
hypre_MultiParVectorMultiScale
hypre_MultiParVectorAxpy
hypre_MultiParVectorMultiAxpy
hypre_MultiParVectorInnerProdGram
hypre_MultiParVectorInnerProdDiag.

The main problem we confronted while creating these functions was the presence of a multivector mask. The vectors are divided by active and passive parts—deflated vectors are passive and others are active. The mask is an array of ones and zeros. The i-th component is equal to one if i-th vector is active; otherwise, it is equal to zero. Because of the presence of a mask, all multivector functions became much more complex.

We changed the communications to reduce the number of messages for computing the inner product. We also changed the communications that are necessary for matrix-vector multiplication. These changes reduced the total number of messages and consequently, the time necessary for computations.

The main users of multivectors will be block iterative methods. One such method is the Locally Optimal Block Preconditioned Conjugate Gradient (LOBPCG) method for solving large eigenvalue problems. Based on general Hypre interface interpreter, we created an interface interpreter that gives us the ability to use abstract pcg code for multivectors.

In collaboration with Jean-Luc Fattebert (CASC), we solved several eigenvalue problems with very large matrices that came from real problems of electronic structure calculation. Coincidentally, we obtained similar results using LOBPCG and preconditioned (with multigrid preconditioner) gradient descent.

As a product of our joint work with Panayot Vassilevski (CASC), a paper titled, "On the Convergence Rate of the Conjugate Gradient Method Exploiting Arbitrary Initial Search Vectors," (P. Vassilevski, A. Knyazev, A. Jujunashvili) has been submitted.

I also collaborated with John Pask (H Division) and his summer student, Kristopher Andersen, to make multigrid preconditioned conjugate gradient methods for a new fortran code intended to solve large eigenvalue problems in electronic structure calculations.

TOC

Automatic Detection of Hot Spots in C/C++ Programs

Student **Alin Jula,** Texas A&M University

Mentor **Dan Quinlan,** CASC

My work at LLNL contributed to the automatic detection of hot spots in C/C++ programs. It contains two separate tools: the *outliner* and the *generic instrumentation tool*.

The outlining mechanism was implemented in the ROSE source-to-source translator, a project in CASC division overseen by Dan Quinlan. The outliner extracts portions of code from a source program into a standalone module and calls to the outlined modules are inserted in the remaining program. The outliner has the inverse functionality of the inliner. If outlining and inlining are performed on a program, the resulting program is semantically equivalent with the initial program. The user inserts special pragma statements that mark the portions of the code that need to be outlined.

The generic instrumentation tool allows programs to be instrumented with *any* instrumentation tools (TAU, PAPI, etc.) using a simple instrumentation format. This format is independent of the tool used, and contains the rules about where and what will be instrumented in the program. The "where" question is answered by using an expression in which the operands are pivotal concepts in a C++ program (loops – for, while, do-while, files, functions, scopes and function calls). The operators are the logical operators (,), and, or, not and in. The evaluation of an expression results in a collection of points in the program for which the instrumentation will be performed. The "what" question is answered by reading the user instrumentation statements from a separate file. This tool is useful for large programs where hand instrumentation could be extremely difficult and error prone.

The generic instrumentation tool is used for empirically detecting the hot spots in a program, while the outliner is used to extract the identified hot spots for further optimization.

Automatic SIDL Generation with ROSE

Student

Kirk Kelsey, University of Rochester

Mentor

Tom Epperly, CASC

The goal of my summer work was to allow for easier use of the Babel language interoperability tool. Babel lets a programmer access a software library implemented in one supported language to be called from any other supported language. Babel supports C, C++, Fortran 77, Fortran 90, Python, and Java. Its language interoperability is based on specifications written in the Scientific Interface Definition Language (SIDL). Typically a programmer would have to write the SIDL code by hand after writing the software library. My goal was to automatically generate the SIDL from the source code of the software library.

The generation of SIDL is being performed using the ROSE preprocessor generator. ROSE is a set of libraries that can be used to parse C source code and generate new code based on the structure of the original program. Although ROSE is typically used to generate more optimized C code, in this case it is being used to generate SIDL.

Once the SIDL is generated by ROSE, Babel must be invoked to generate skeleton code in each supported language to wrap the original library. This wrapping code facilitates access from other languages. Currently, the skeletal wrappers must be completed by the writer. The next stage of this work is to use a ROSE preprocessor to complete the skeleton code automatically. Once this is done, language interoperability will be achieved without any writing needed beyond the original library.

Large-Scale Atom Data Visualization

Student
Jason Kimball, University of California, Irvine

Mentor
Mark Duchaineau, CASC

Scientific visualization of the large sets of atoms produced by numerical physics simulations at the Lab presents specific challenges due to the large number and geometric nature of these atoms. Rendering data sets on the order of 10–100 million atoms requires the use of parallel rendering clusters that may not be able to provide interactive rendering speeds (60 frames per second with minimal latency), especially when the user's computer has a slower connection to the rendering cluster. Also, since atoms should be rendered as unique spheres, traditional, performance-improving, level-of-detail (LOD) algorithms based on geometry simplification are useless because simplifying the geometry of a large group of atoms removes the granularity of each atom.

Building upon the ROAM architecture developed by Mark Duchaineau, I implemented a perspective-based LOD system that renders progressive-quality views of the data. Each rendering becomes a texture on the wall of a cube that surrounds the "camera" viewing the data. From within this cube, the user is free to look in any direction and see data, at least in a low-resolution format. Meanwhile, the system renders progressively higher resolution views for the interesting regions of the data. Being able to request higher-resolution renderings for only limited areas of interest greatly reduces the rendering work. Furthermore, since the user has a view of the data at all times, this technique provides the desired interactive speeds, and the progressive resolution updates provide detailed information about the data.

In addition to overcoming the aforementioned problems with atom data, this solution works for data with very complex geometry and therefore, does not work well with traditional LOD algorithms.

Future work includes developing algorithms to optimize the selection of "interesting" regions, as well as using several different viewpoints and image warping to provide movable cameras.

TOC

Algebraic Multigrid for Maxwell's Equations with Variable Coefficients

Student

Nicholas Kridler, University of Colorado at Boulder

Mentor

Barry Lee, CASC

There are many difficulties in creating an efficient multigrid method for the curl-curl formulation of Maxwell's Equations. The key difficulty is the large near-null space of the curl operator. While there are several successful algebraic and geometric multigrid methods for this problem, performance degrades when there is a large variation in the material coefficients. The reason for this lies in the interpolation operator. In the approach by Hiptmair [3], the interpolation was constructed under the constraint that a multilevel commutativity complex is formed. However, under these constraints, the interpolation operators cannot handle divergence-free error components.

Ideally, one would like a complex that includes the divergence operator, but such an operator is extremely difficult to construct. Lee and Jones [4] constructed a scheme that relaxed on the multilevel commutativity constraint, allowing more freedom in the construction of the interpolation operators to handle both curl-free and divergence-free error components. In their approach, the null space of the curl operator is explicitly introduced into the set of equations. For this formulation, interpolation must be constructed for the edge and nodal elements. The nodal interpolation is constructed by box-mg [2] and the edge interpolation is constructed by AMGe [1].

Constructing the edge interpolation requires defining the weights for the boundary edges on the fine level and then solving a local problem for the weights corresponding to the interior. The ratio of the length of the fine edge to the length of the coarse edge is used to define the boundary weights. Unfortunately, this does not take the variable coefficients into consideration and will fail if there is too much variation. My project involved investigating ways to create edge interpolation based on the variable coefficients when only given a user-defined fine-grid matrix.

In Maxwell's Equations, the coefficient of interest is the magnetic permeability. Magnetic permeability is related to magnetic flux through a constitutive relation. Our goal was to create interpolation based on enforcing tangential continuity of the magnetic flux through the jump discontinuity of the coefficient. This requires assuming that all of the boundary weights are unknowns, and it is not clear that the problem will remain local. It is quite possible the fine edge I am trying to construct depends heavily on a neighboring element. The other problem is that the magnetic permeability is unknown, so it may not even be possible to enforce the constraint.

Another approach involves ignoring the fact that the coefficient is the magnetic permeability and looking at it as a stretching or shrinking factor. The effect of the coefficient can be seen as the transformation of the grid and weights could again be chosen by ratios of lengths. It is not clear that this transformation gives us any information we can use or if there are too many unknowns for it to be effective. Both approaches are still under consideration. Ideally, we would like a method that uses only values along a single coarse edge to create the corresponding fine grid elements.

References

[1] M. Brezina, A. J. Cleary, R. D. Falgout, V. E. Henson, J. E. Jones, T. A. Manteuffel, S. F. McCormick, J. W. Ruge, Algebraic multigrid based on element interpolation (AMGe), SIAM J. Sci. Comp., 22 (2000), 1570-1592.

[2] J. E. Dendy, Jr., Black box multigrid, J. Comput. Phys., 48 (1982), 366-386.

[3] R. Hiptmair, Multigrid methods for Maxwell's Equations, SIAM J. Numer. Anal., 36 (1999), 204-225

[4] J. E. Jones and B. Lee, A Multigrid Method for Variable Coefficient Maxwell's Equations, SIAM J. Sci. Comp., submitted.

TOC

IMAGE Pipeline

Student **Brian Lum,** University of California, Berkeley

Mentor **Amber Marsh,** EEBI Division

I worked on a group project called the Integrated Molecular Analysis of Genomes and their Expression (IMAGE) Pipeline Project. I helped the group meet deadlines by taking on a variety of tasks that otherwise would have had to become post-production projects or would have forced the deadline to be pushed back.

My first task was setting up barcode printers. The printers required a great deal of complicated configuration, and it was my job to work with the printer company and with the Laboratory's information technology (IT) group to ensure that everything was set up to meet our needs. I also determined how to program the printers and how to design a printer command generation module that integrated with the rest of the Laboratory's Web-based interfaces.

My second task was designing and implementing a Web-based interface to maintain the personnel database related to the project. I also retrieved and formatted the data that I eventually used to populate the initial personnel database.

In a similar side project for IMAGE, I worked with the project's Principal Investigator (PI), Christa Prange, to retrieve outdated information from an old database, update it, reformat it, and enter it into the new system. I then designed and implemented a Perl script that is run nightly via CRON to display the information.

Numerical Solution Methods and Error Analysis for the Neutron Transport Equation in Absorbing, Monoenergetic, Non-Scattering Media

Student **Eric Machorro,** University of Washington

Mentor **Britton Chang,** CASC

We considered the local truncation errors and estimated global errors for the exponential-characteristic (EC) method for solving the 2D monodirectional, non-scattering, Boltzmann neutron-transport equation. We showed that under ideal circumstances, the local error is second order, yet globally, the method is only first-order accurate. Test problems were evaluated, and in some respects, the EC method performed better on problems with discontinuous material interfaces. An analytic counter-example was explored demonstrating at best O(h) accuracy.

As an alternative, we developed the Boundary Exponential Log Characteristic (BELCh) method. The method is positive and locally satisfies the conservative form of the equation. Estimates for local truncation errors and global errors for the method are under development.

For strictly positive solutions that have L-continuous derivatives, we show that BELCh is locally Lth order accurate, yet globally, the method is only L–1-order accurate. Examples and test problems were explored, and in some important respects, BELCh outperforms the EC method.

We also proved that the Petrov-Galerkin discretization of the monodirectional neutron transport equation is convergent. This was demonstrated by proving that the inverse of the matrix for the system is bounded by a number that is independent of the order of the matrix. Therefore, the global error of the Petrov-Galerkin approximations of the transport equation has the same order as the local truncation error.

Out-of-Core Visualization of Climate Modeling Data

Student **Anna Majkowska,** University of California, Riverside

Mentor **Valerio Pascucci,** CASC

My research activity during last summer has been focused on enhancing the ViSUS IDX streaming infrastructure. The work has been centered around two main tasks:

- Extension of the data extractor in the streaming infrastructure to allow synchronized retrieval of multiple scalar fields.
- Development of a new out-of-core isosurface extraction algorithm that takes advantage of hierarchical Z-order data layout. The scheme allows for view dependent refinement and avoids storing explicit data structure for the triangles in the isosurface. In particular, the vertices of the isosurface are computed by a vertex program that takes a tetrahedron T and then renders the portion of isosurface contained in T.

The main test bed for the new features has been the 3D image data obtained from the IMPACT climate modeling simulation. With the current enhancement, one can connect the ViSUS streaming framework directly to the simulation, and while the simulation is running, visualize isosurfaces of the ozone concentration pseudocolored by temperature distribution (see figure below). The synchronized streaming allows for the introduction of streaming computation of derived scalar fields, which is a crucial component of any complex visualization environment for large scientific data.

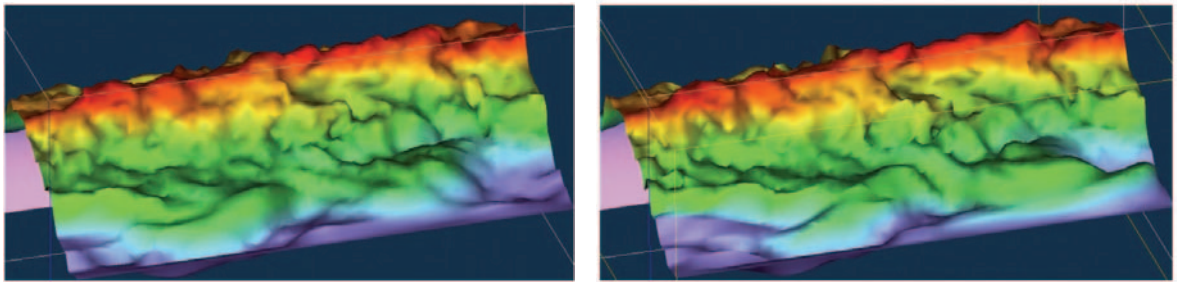

Two time steps of the IMPACT climate modeling simulation visualized with an isosurface of the ozone concentration pseudocolored with second scalar field.

TOC

Code-Independent Analysis Tools for Physics Simulation Codes

Student **Robert Dean Malmgren,** Northwestern University

Mentor **Scott Brandon,** AX Division

It is becoming increasingly important for computational physicists to be able to analyze the accuracy of their codes during calculation as well as in postprocessing. While visualization techniques are readily available, simulation precision cannot be judged visually. Furthermore, it is impractical to visualize the problem during computation to judge accuracy. Code-independent analysis methods, which offer both in-line and postprocessing capability, are therefore developed to give computational physicists a means of rigorously analyzing any code's performance.

The developed tool quantitatively measures directional symmetry, relative error contours, and relative shock front location based on the computed and analytic solutions. Inline implementation of this tool gives computational physicists the ability to have their codes adapt and self assess throughout, for example, a parameter optimization experiment. In postprocessing mode, these techniques allow computational physicists to examine error patterns in their codes and ultimately facilitate code development and fine tuning.

The following figures illustrate the tool's capability compared to KULL's performance against the predicted analytic solution of the Noh Problem. Figure 1 demonstrates the ability of the symmetry analysis tool. Figure 2 shows the shock front tool's use. Both tools can be used to quantitatively describe asymmetry when the asymmetric profiles are chosen. In this case, the profiles along the axes are asymmetric.

Future work on this project could extend the tool's capability to include an asymmetry finder. Additionally, this tool could be altered to quantitatively compare code performance against experimental data or another code's performance.

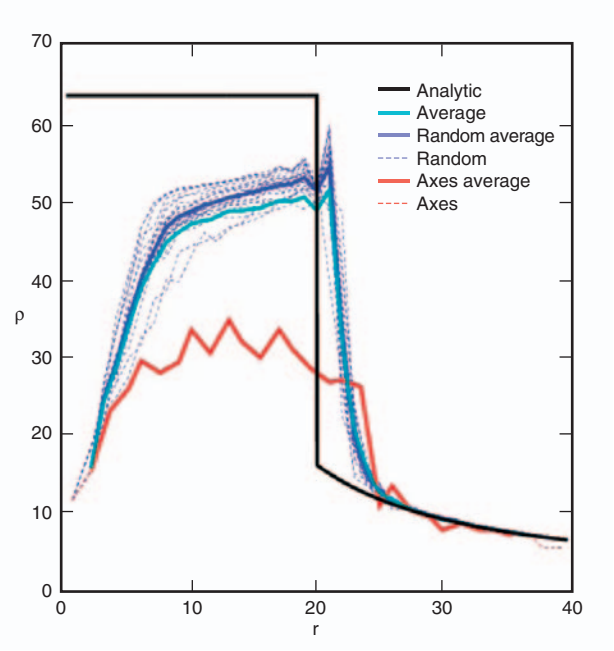

Figure 1. Symmetry analysis of KULL solution to Noh problem

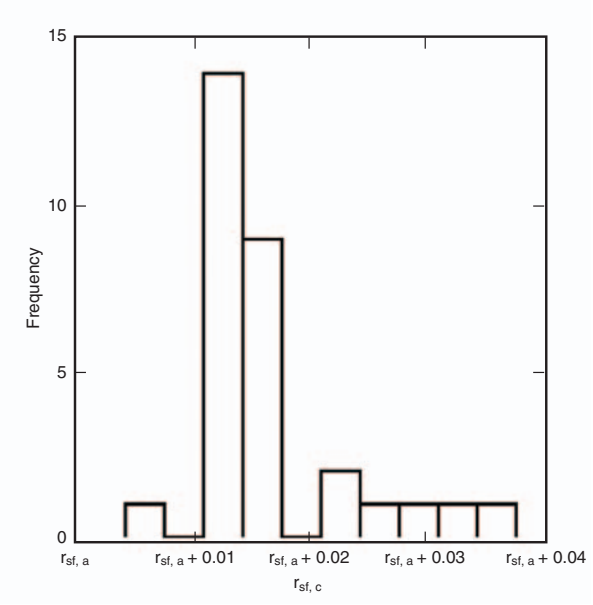

Figure 2. Histogram of relative shock front position of calculated profiles

Development of a Graphical User Interface for the Virtual Beam Line (VBL) Simulation Software

Student

Peter Manning, University of Arizona

Mentor

Glenn Goderre, NIFE Division

The purpose of the project was to develop a graphical user interface (GUI) to interact with and utilize the already-developed VBL software. Major accomplishments include: a method for viewing and editing part properties (parts of the NIF laser); the ability to open any tabulated aberration file, gain file, or beam dump file and view its distribution in 2D and 3D; a mechanism for inserting beam dumps anywhere along the path of the beam and also the ability to remove them; the ability to search for a part and have its information displayed; the ability to recall previous runs and their summaries. The GUI displays the output summary from a VBL run in a readable and configurable format.

I learned more Java techniques and design patterns than I thought possible in a single summer. After 10,000 lines of code and 160 classes, there exists a program that enables users to easily work with the VBL simulation software. Although the GUI will change in many ways in the next few months, these accomplishments provide the foundation for future changes and advancements.

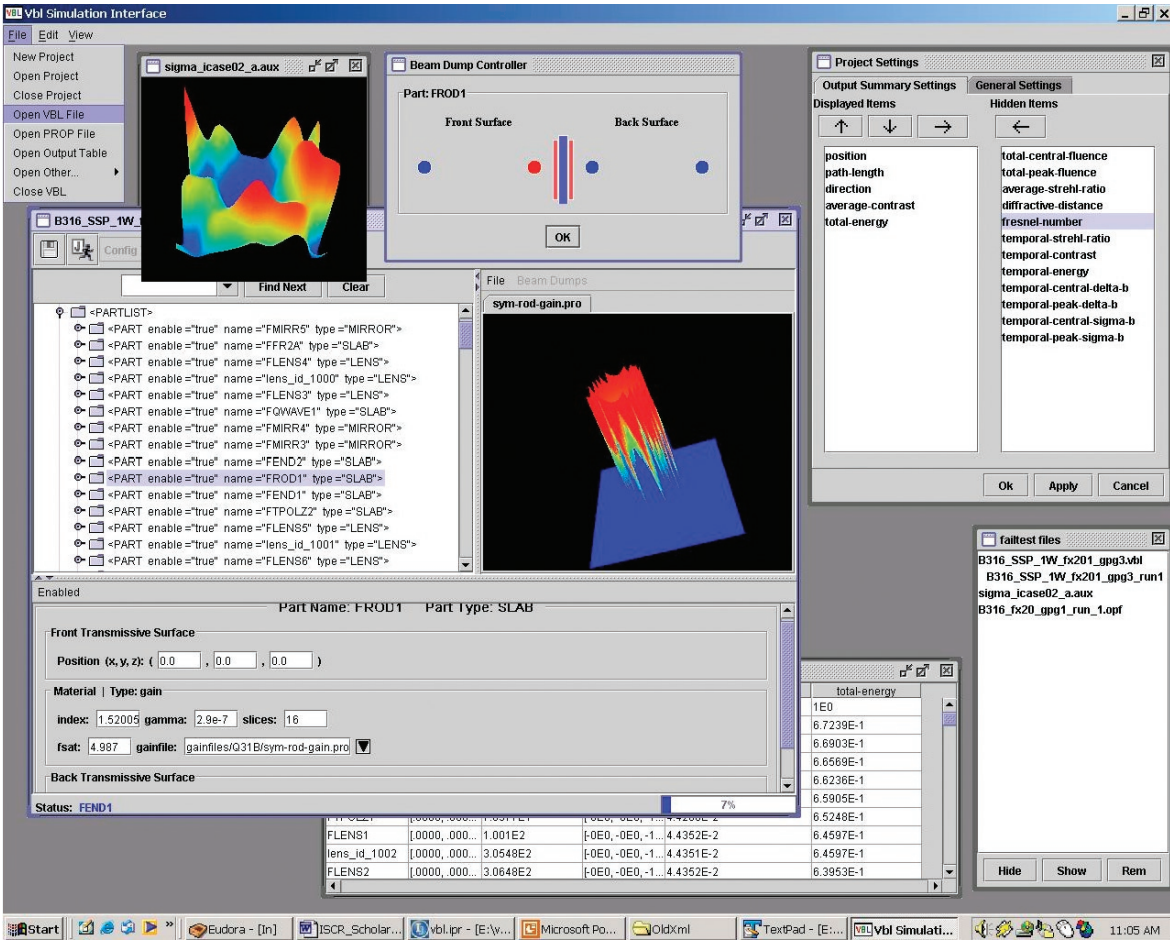

Figure 1. Screen shot of VBL GUI and most of its capabilities during a simulation run

Time-Varying Reeb Graphs for Space-Time Data

Student
Ajith Mascarenhas, University of North Carolina at Chapel Hill

Mentor
Valerio Pascucci, CASC

Physical processes that are measured over time, or modeled and simulated on a computer, can produce large amounts of data that must be interpreted with the assistance of computational tools. Such data arises in a wide variety of studies, including computational fluid dynamics, oceanography, and climate modeling. The data typically consists of finitely many points in space-time and a measured value for each.

Graphical visualization, often through level sets or isosurfaces of a continuous function, is useful for interpreting the data. A level set consists of all points in the domain whose function values are equal to a chosen real number 's'. In three dimensions, this is generically a surface that is then displayed. By varying 's', we can study the variation in the data. Topological features of the level sets, such as connected components, handles, and voids, can be important aids in interpreting the data. By encoding the evolution of these features, the Reeb graph compactly represents topological information of all level sets. As we pass through time, the Reeb graph goes through an evolution of its own, undergoing structural changes at birth-death points and at interchanges of critical points. The evolution of the Reeb graph thus represents a 2-parameter family of level sets. We suggest that this 2-parameter family, encoded in a compact data structure, is a useful representation of space-time data.

I developed software that computes time-varying Reeb graphs specialized for the case of 2D scalar fields. I first worked on the mathematical foundations of Morse functions and then addressed the practical issues that allow a robust implementation. A central challenge in this work is defining a set of consistent rules to handle the degeneracies present in real-world data. These policies bridge the gap between the theory for smooth functions and the practical world of sampled fields obtained from simulations or experiments. The software development is in progress and will be finalized in the next two months at the University of North Carolina at Chapel Hill.

I also started the next stage of this research, involving the construction of multiresolution Jacobi sets, which are the 1-manifolds that describe the trajectories of critical points over time. I implemented the initial ideas in a software prototype that visualize the evolution of a 2D scalar field. In this tool, the Jacobi sets are first presented at full resolution and are then incrementally simplified to highlight the most important features in the data. The preliminary tests with combustion data sets show promising results. The figure below shows the Jacobi set of the combustion data with three levels of simplification, starting from full resolution at left.

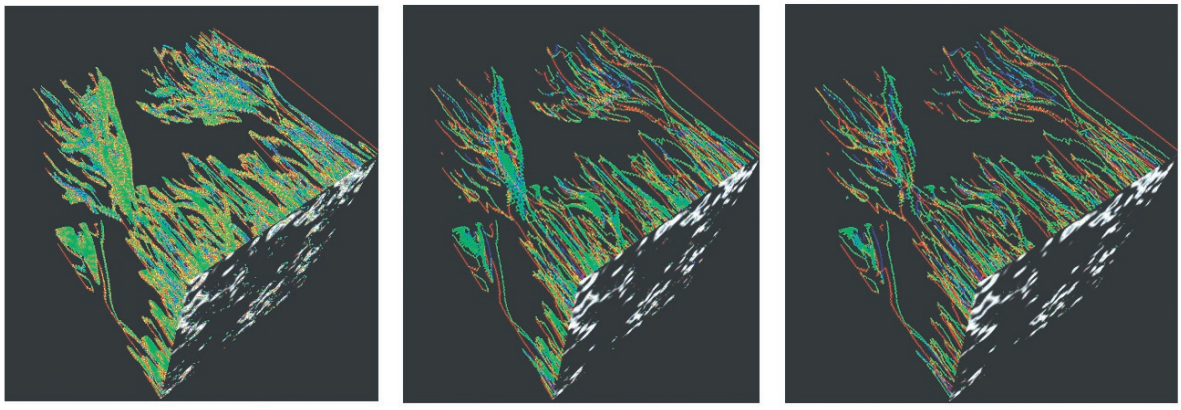

Figure 1. Three levels of simplification starting with full resolution (left) depict a Jacobi set.

Ultra-Wideband (UWB) Token Ring Simulation & Security Needs of Sensor Networks at LLNL

Student **Tammara Massey,** Georgia Tech / University of California, Los Angeles

Mentor **Terry Brugger,** NAIC

Ultra-wideband is an emerging wireless technology that permits covert and reliable transfer of data. It is essential to investigating networking protocols, such as token ring, that allow for quality of service and bounded latency for ultra-wideband radios. To explore the limitations of the ultra-wideband token ring protocol, throughput was measured with varying levels of mobility and with various quantities of nodes. With a larger number of nodes, the throughput increased because each node could transmit data to more possible nodes. The throughput for nodes with mobile topologies was lower than for nodes with stationary topologies due to the additional overhead of mobile nodes joining and leaving rings. The simulation also showed how a larger distance between the nodes and a larger topological area decreased the throughput. The simulation proved that the token ring protocol provided the desired throughput and quality of service for ultra-wideband radios transmitting multimedia. I created a poster and a technical paper describing my research and suggesting further work. With the aid of my mentor, this paper will also be submitted to a conference.

Security in Sensor Networks

Several wireless sensor technologies exist at LLNL that gather data for environmental or defense applications. Unfortunately, wireless technologies are more vulnerable to security attacks because the data is broadcast. The ubiquitous nature of broadcasting allows adversaries to eavesdrop or inject messages into the wireless network more easily than in a traditional wired network. After describing the current sensor networks at LLNL, I analyzed the security vulnerabilities of sensor networks. I presented security solutions that need to be implemented in current and upcoming projects. I also proposed future areas of research that need to be explored to ensure that LLNL sensor applications remain safe and reliable. I wrote a technical paper describing this research.

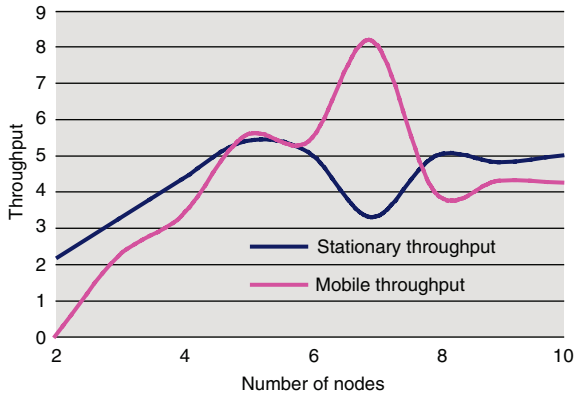

Throughput for Ultra-Wideband Token Ring Simulation with an Increasing Number of Nodes

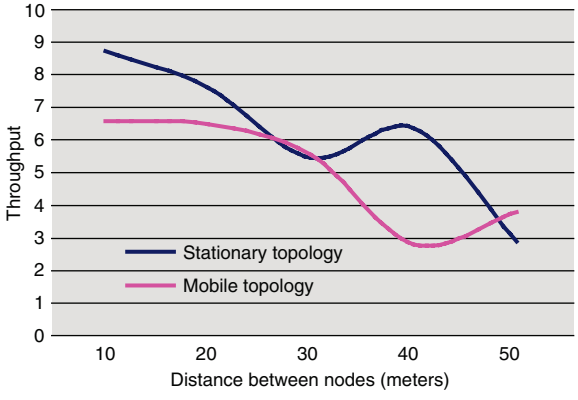

Throughput for Ultra-Wideband Token Ring Simulation with an Increasing Distance Between Nodes

PerfTrack

Student **Kathryn Mohror,** Portland State University

Mentor **John May,** CASC

PerfTrack is a tool to help scientific programmers answer difficult questions about application performance, given that the source code, build parameters, runtime environment, and hardware vary over time. PerfTrack was developed to explore technologies in parallel performance measurement, modeling, analysis, and prediction. Performance data and the associated environment data is being stored in a relational database. This database provides a foundation to build analysis tools that are scalable to large numbers of threads and are capable of comparing multiple executions. Our goal is an automated tool that will gather, store, and analyze data, in order to encourage their use in the software development cycle.

This project was motivated by needs in the high-performance computing community. Code teams need to have an organized way to understand how changes in hardware, build and runtime environments, and source code affect application performance. Machine purchasing teams need a way to predict how future machines will perform so that they can select machines with characteristics that are best suited to the codes that are important to them. PerfTrack addresses these needs by researching solutions for problems that prohibited such tools in the past, such as how to compare possibly disparate performance data, and given a large amount of performance data, how to present only relevant information to the PerfTrack user. In future work, we plan to investigate other research problems, such as how to provide automatic correlation of performance trends with machine, software, and environmental characteristics.

Summarizing Network Traffic with Information Visualization

Student

Christopher Muelder, University of California, Davis

Mentor

Marvin Christensen, NAIC

Lawrence Livermore National Laboratory and several other Department of Energy sites have sensors on their networks that collect vast amounts of data. The quantity of data collected is too large to go through by hand, and automated analysis is not yet completely reliable. A program was developed to display a high-level visual summary so that suspicious behavioral patterns could be easily and rapidly detected. I worked on porting the program to Mac OS X, changing the layout to be more intuitive, improving existing capabilities, adding new capabilities, and fixing many small internal issues. The process of porting the program to OS X gave me the opportunity to get accustomed to the operating system and the existing code, and to learn how programs have to be changed to work on OS X. The layout of the program was rather unintuitive, flowing in essentially a counterclockwise direction. The new layout flows from left to right and is smaller overall so that it will work on screens as small as 1024x768. One of the visualizations was labeled in order to make it easier to use, and the histograms were improved to increase their color association and make it easier to edit custom gradients. New capabilities include the ability to normalize the data so that data from different parts of the day can be compared, and the ability to calculate variance relative to the values averaged across time periods, as well as an absolute variance.

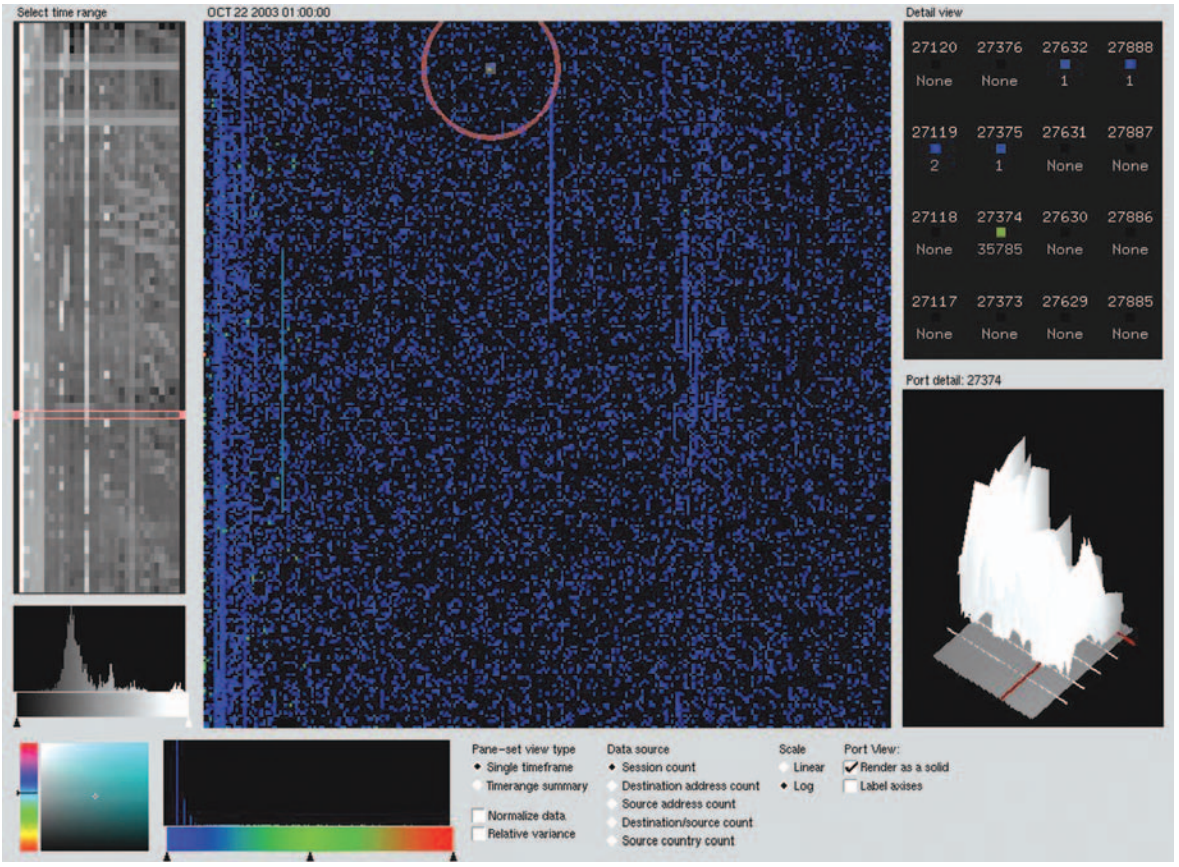

PortVis is a graphical interface that allows computer security analysts to quickly sort thought large amounts of connection data to identify interesting abnormalities that warrant further investigation. Using this tool is a quick way to filter out normal network activity and focus on suspicious activity regardless of the variance of malicious code.

TOC

Software Failure Risk Analysis

Student **Spencer Nielsen,** Brigham Young University

Mentor **Darrel Whitney,** CADSE

My assignment was to create a Web-based application to enable project managers at the Lab to grade and assess their projects for likelihood of software failure and the risk consequences associated with failure. After the likelihood of failure and risk consequences of a project have been reported, the project is assigned a calculated grade. This grade is then used to consider practices and policies to implement in the project and to make other management decisions. I also wrote an administrative interface to allow directorate-level management to customize the evaluation process for projects in their directorates. I used the PHP/MySQL development platform to accomplish this task.

A Space-Time Sparse Grid Approximation Space

Student **Daniel Oeltz,** Rheinische Friedrich-Wilhems Universität Bonn

Mentor **Panayot Vassilevski,** CASC

In this project, we developed and implemented new approximation spaces for parabolic problems, so-called space-time sparse grid spaces. With classical space-time discretization spaces, using a uniform grid in d-dimensional space with N-grid points in each coordinate direction usually needs $O(N^{-(d+1)})$ or $O(N^{-(d+2)})$ degrees of freedom. Those spaces will reduce the number of degrees of freedom to $O(N^{-d})$ or $O(N^{-(d+1)})$. In the first step, we analyzed the approximation rates of those spaces. We were able to prove for special function classes that, besides the tremendous reduction of the degrees of freedom, the space-time sparse grid spaces provide nearly the same approximation rate as classical space-time Finite Element spaces.

In opposition to classical sparse grid approaches, we are using a tensor product construction of a d-dimensional multilevel basis in space and a one-dimensional multilevel basis in time, which makes it possible to handle even complicated domains in a very easy manner. Here, we construct the spatial multilevel basis from a spatial Finite Element discretization via AMG coarsening and a least square approach. We implemented this approach, including adaptivity, and performed several numerical experiments showing that this approach is able to approximate space-time functions with great accuracy while reducing the number of the overall degrees of freedom significantly compared to Finite Element spaces. It turned out that by applying adaptivity, this approach provides good results, even for highly singular functions.

Reviewing Vista

Student

Christopher Olson, University of California, Santa Cruz

Mentor

Robert Cooper, DCOM

Computation loops in ALE3D involve a fixed notion of an index map. I implemented two easy methods to abstract this indirection in an effort to reduce table lookups and better respect cache lines. Unfortunately, the cost of both techniques, especially from register spilling due to additional register pressure, outweighed the anticipated benefits. Other techniques may still allow us to clarify code and reduce the runtime of computation loops, but they will require more effort to build, and they may involve significant changes to the existing source code. I also rewrote ALE3D's Vista library to get a cleaner and faster implementation.

Subspace Detectors

Student

Timothy Paik, University of California, Berkeley

Mentors

Shawn Larsen & Dave Harris, ISCR

When handling seismic signals, it is often necessary to correlate different signals against each other. One of the tools used to correlate seismic signals is the subspace detector. Subspace detectors calculate template signals from events and correlate them against other events. Subspace detectors are useful because they return correlation values independent of the amplitude of the signal. This behavior is helpful for two reasons. First, they do not return misleading large correlation values when receiving large signals. Second, they can find correlating seismic events that might be buried under the noise of other unimportant signals.

My assignment this summer was to write Java classes designed to improve the signal-processing engine of the subspace detector. I wrote several classes to represent complex signals and several other classes to process complex signals. My main goal was to improve the overall speed of the subspace detector. I achieved this goal through a change in the algorithm and the development of a quick Fourier transform. The subspace detector now runs about 3.5 times faster than before, allowing for greater amounts of subspace detectors to be deployed in the field.

TOC

Streaming Pointsets

Student

Sung Park, University of California, Davis

Mentor

Peter Lindstrom, CASC-ISCR

The increasing amount of point data generated by different applications is posing a major challenge in data processing. Typically, such large data sets not only exhaust the main memory resources of common desktop PCs but can even exceed the 4 GB address-space limit of a 32-bit machine. To process data sets that do not fit in main memory, out-of-core (external-memory) algorithms are used.

Dealing with large data sets for the purpose of scientific visualization is an ongoing area of research. For large point data, there has not been much work to address this problem. One ideal approach for dealing with large point data is a streaming model. In a streaming model, data is streamed into the main memory and retired as soon as it is finished processing. My supervisor's previous work has shown that this approach works well for large meshes.

My work this summer has been trying to come up with an optimal approach to do stream processing on point-based data sets. The problem proved to be a lot harder than anticipated as many different approaches we tried gave poor results. The problem is difficult because pointsets do not provide connectivity information, thus finding even neighboring information requires certain assumptions about the data that has been streamed into the machine. The approach I recently investigated has shown some good results and deals with finalizing partitions of space by using a binary partition tree. The work is still in its beginning stages but preliminary results have been good, and we are hoping to further investigate this idea.

Internet Ballistics: Retrieving Forensic Data From Network Scans

Student **Bryan Parno,** Harvard University

Mentor **Tony Bartoletti,** Computer Incident Advisory Capability (CIAC)

The typical network receives millions of hostile probes every day, and a significant portion of these probes are network scans. During a network scan, the attacker sends connection requests to every possible network address and listens for replies indicating the presence of a (possibly vulnerable) computer. Since a network scan often serves as a precursor to an attack, reliable identification of scanners can significantly enhance cyber-security. Furthermore, the ability to map adversary hierarchies and correlate the attacks with events in the real world contributes to counterintelligence work.

For a variety of reasons, source IP addresses fail to provide the necessary identification information. However, analyzing packet arrival timing data reveals highly distinctive patterns that may correlate with the attacker's choice of tools, physical platform and/or network location (Figure 1). By selecting data transforms conducive to periodic analysis, we can use wavelet techniques to achieve more than 1,000x compression ratio while still preserving the essential features (Figure 2). Initial experiments indicate our methods consistently identify patterns in the data. In future work, we plan to perform controlled scans using common network scanning tools from multiple locations to refine our identification techniques, allowing us to reliably fingerprint network scanners, without relying on the source IP address.

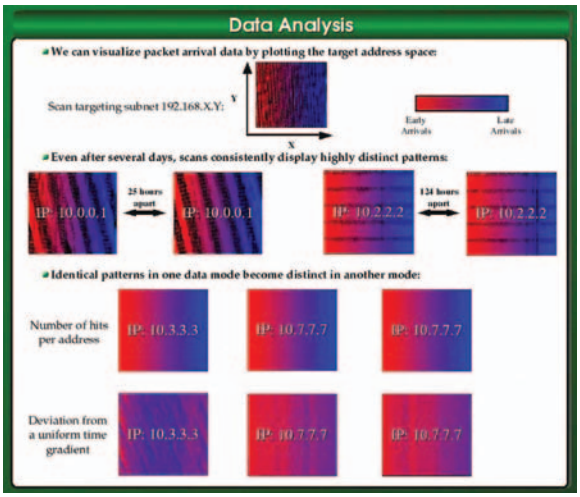

Figure 1. Class-B subnet plots of various hostile packet statistics (packet arrival times, hits per target address, and others) reveal multiscale structure often characteristic of source."

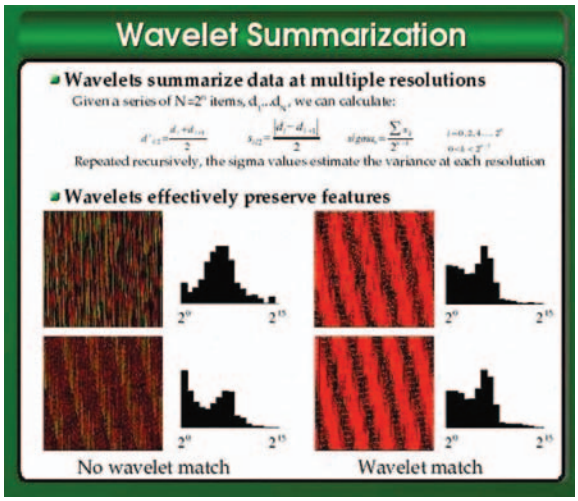

Figure 2. Wavelet scalograms reliably capture and record multiscale structure in scan-traffic statistics, providing compact "scan-behavior signatures" for identification purposes.

SLURM for BlueGene/L

Student **Dan Phung,** Columbia University

Mentor **Moe Jette,** CASC

My duties this summer were related to the incoming system installation of IBM's BlueGene/L (BGL), currently the world's fastest computer. I was given the task of porting the code that currently manages the resources of the other large cluster systems for the BGL system. My other responsibility was to design and implement a partitioning and wiring algorithm that divides the system into portions to allow the system to be shared among users and efficiently utilizes the system.

Software Vulnerability Taxonomy Consolidation

Student **Sriram Polepeddi,** Carnegie Mellon University

Mentor **Noel Tijerino,** IOAC

The goal of the software vulnerability taxonomy consolidation project is to address the need for a universally accepted vulnerability taxonomy that classifies vulnerabilities in an unambiguous manner. In today's environment, computers and networks are increasingly exposed to a number of software vulnerabilities. Information about these vulnerabilities is collected and disseminated via various large publicly available databases, such as BugTraq, OSVDB, ICAT. These databases do not cover all aspects of a particular vulnerability and lack a standard format, making it difficult for end-users and administrators to easily combine and compare various vulnerabilities. A central database of vulnerabilities has not been available until today for a number of reasons, including non-uniform methods by which current vulnerability database providers receive information, disagreement over which features of a particular vulnerability are important and how best to present them, and the nonutility of the information presented in many databases.

A consolidated vulnerability database (CVDB) was implemented that clearly communicates information between the vulnerability discovery agent and end users of that data. Data from BugTraq, OSVDB, ICAT, Secunia, Computer Associates' Vulnerability Information Center, and Mitre's CVE was fed into a database. These sources required extracting information directly from Web pages, XML, CSV and Tab-delimited files. Once loaded into a database, the data most unambiguously representing an aspect of a vulnerability was loaded into the main CVDB tables and linked using CVE IDs. While representing a smaller subset of the total available data, the information now possesses a precise and repeatable structure.

TOC

Performance-Oriented, Privacy-Preserving Data Integration

Student
Raymond Pon, University of California, Los Angeles

Mentor
Terence Critchlow, CASC

Data is often generated or collected by various parties, and the need to integrate the resulting disparate data sources has been addressed by the research community. Although the heterogeneity of the schemas has been addressed, most data integration approaches have not yet efficiently addressed the privacy requirements imposed by data sources. Without strong privacy guarantees, scientists often refuse to share data with other scientists for reasons such as subject/patient confidentiality, proprietary/sensitive data restrictions, competition, and potential conflict and disagreement.

When sharing scientific data, privacy quickly becomes an issue. Suppose that a scientist wishes to perform a query across a table in his private database and a table in a public data warehouse in the most efficient manner possible. Ignoring privacy restrictions, the problem is reduced to a distributed database problem that can be solved by shipping the scientist's table to the data warehouse and performing the join at the warehouse. However, if the scientist's data set is proprietary, it cannot be sent verbatim to the data warehouse. The naive solution is for the scientist to download the entire data warehouse to his local machine and perform the query there. But to do so would be prohibitively expensive if the data warehouse is very large or the communications link is limited.

To address this problem, we augment the well-known semi-join framework, "hiding" the actual values of the join column of table *R* by hashing them and including additional artificial values. Our hash/noise method approach uses a set of fixed hashing and fake hash values to control the amount of uncertainty in the identity of the join column values in table *R*, thereby controlling the level of privacy loss incurred. By sacrificing a small fraction of privacy, this method incurs significantly less transmission costs than the naive solution. Relative information gain is used as the measure of privacy loss.

Experiments were done to evaluate the effectiveness of this approach. The results of the experiments will be published at the next appropriate conference.

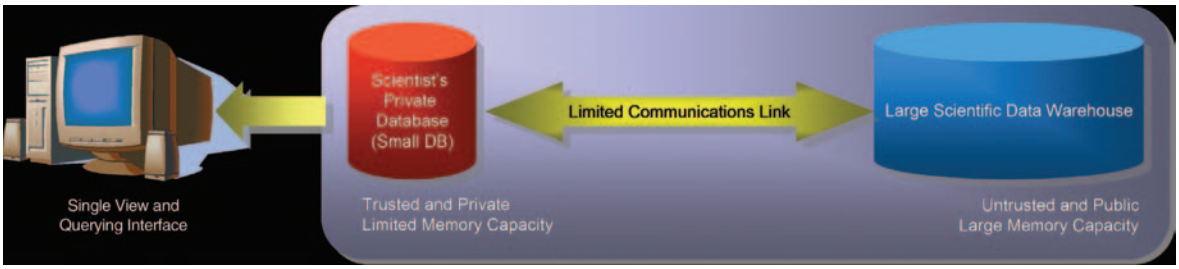

Target architecture for privacy preserving queries

Out-of-Core Parameterization and Remeshing of Surfaces

Student **Serban Porumbescu,** University of California, Davis

Mentor **Mark Duchaineau,** CASC

The goal of this project is to perform remeshing and parameterization of surfaces too large to fit into the main memory of conventional computers.

Our approach creates a hierarchy of fine-to-coarse resolution surfaces via repeated removal of an independent set of vertices. The independent set consists of non-neighboring vertices. This construction process results in a hierarchy of topologically equivalent surfaces and a robust one-to-one and onto mapping between these surfaces. Consequently, a point on a surface within the hierarchy can be mapped uniquely to any other surface in the hierarchy.

With this hierarchy at our disposal, we perform remeshing by regularly subdividing the coarsest level of the hierarchy and mapping the newly inserted vertices to a higher-resolution surface within the hierarchy. This mapping is generally not smooth and the introduced elements are highly distorted. To ameliorate this problem, we apply a smoothing filter to our remesh surface that regularizes elements with respect to a particular level of the hierarchy.

We are currently exploring several approaches to smoothing of the remesh surface with respect to a surface in the hierarchy. The first approach involves smoothing the remesh surface with respect to a portion of the target surface embedded in the plane. An alternative approach involves smoothing with respect to a locally constructed scalar field.

Future work involves mapping our algorithms and data structures to their out-of-core counterparts. While not trivial, our highly localized algorithms and data structures should help facilitate a smooth transition to an out-of-core approach.

TOC

Intrinsic Function Testing and an OpenMP Runtime Library

Student

Ashley President, Carnegie Mellon University

Mentor

Mike Kumbera, DCOM

The purpose of the project is to add necessary and useful features and to test existing and new features of GNU gfortran, an open source Fortran 95 project. The feature extensions that needed to be added are OpenMP, Application Program Interface (API), and Cray Pointers. Along with the added features, the current features of Fortran 95 need to be thoroughly tested.

Some of the intrinsic functions, required by Fortran 95 specifications, needed to be tested in gfortran. Once the tests were created, they were used to check the functionality of several preexisting Fortran 90/95 compilers. Forty-eight different functions were tested.

OpenMP is an Application Program Interface (API) that can direct multi-threaded, shared memory parallelism in Fortran, C and C++. OpenMP calls work with a set of simple directives that explicitly control the programs parallel activities. Along with directives, the OpenMP standard also specifies a set of runtime library calls that can be used to determine the system state. The OpenMP Runtime Library is designed to work with both C/C++ programs as well as Fortran Programs. The backbone of the runtime library is written in C and uses POSIX (pthreads)to implement threads.

While many of the gfortran features were tested, many other features within the compiler still need testing. The OpenMP runtime library stills needs to be tested on a larger scale. The C/C++ compilers need to be updated to recognize OpenMP calls.

TOC

AS2TS Web site

Student **Davinder Rama,** California State University, Sacramento.

Mentor **Adam Zemla,** EEBI

The Amino-acid Sequence into Tertiary Structure (AS2TS) and Local Global Alignment (LGA) systems provide a capability to create 3D models of analyzed proteins, detected signature regions and other biologically important fragments of protein structures. The goal of my project was to develop a publicly accessible Web site that will allow submission of queries to the selected local services:

Protein Structure Analysis services (LGA):

- Protein structure comparison facility — Finds 3D similarities in protein structures. Allows the submission of two 3D protein structures for structure comparative analysis.
- LGA–AS2TS Models structure comparison — Allows the submission of the protein structure in Protein Data Bank (PDB) format for LGA comparison with models generated by AS2TS system.
- LGA–PDB chains structure comparison — Allows the submission of the protein structure in PDB format for LGA comparison with chains from the PDB list.

Protein Structure Modeling services (AS2TS):

- Sequence-structure Alignment into Tertiary Structure — Generates a tertiary structure (3D model) for a given sequence-structure alignment model.
- AS2TS Service — designed to facilitate the modeling of tertiary structure (3D model) for a given sequence of amino acids.
- AS2TS Model Builder — An expanded version of AS2TS that allows submission of a set of protein sequences for high throughput protein structure modeling.
- Local Sequence Library for AS2TS system — Create and deposit Libraries in FASTA formatted protein sequences or combine existing libraries.
- Local Structure Library for AS2TS system — Deposit PDB structures in PDB format into the local PDB.
- SCWRL (Side Chain placement With a Rotamer Library) — side chain builder for AS2TS system. For a given protein structure SCWRL calculates new conformation of side chain atoms.

This project required establishing efficient data exchange (query<->results) between the web-server and the computation servers. An overview of the system follows:

The AS2TS website: http://as2ts.llnl.gov

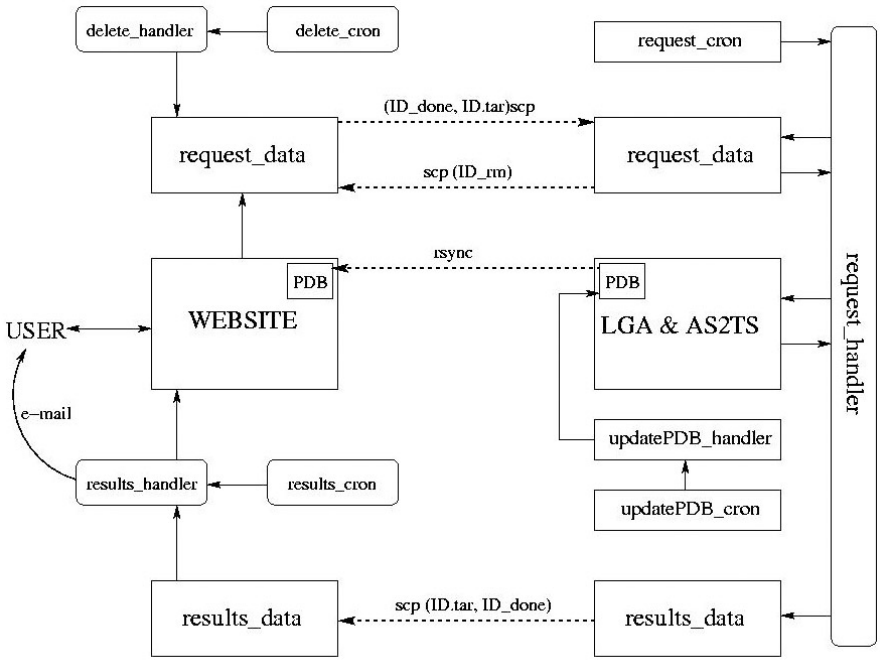

The figure diagrams the interfaces and the data flow between the major modules of the developed system for data exchange between the "WEBSITE" server and the computation "LGA & AS2TS" server. To process user queries, three types of modules were developed—request_handler, results_handler, and delete_handler. Each module is invoked by a cron, and the data exchange is done using scp (secure copy, a remote file copy program).

TOC

Variance-Based Feature Tracking

Student **George Roberts,** Georgia Institute of Technology

Mentor **Chandrika Kamath,** CASC

Many modern computer vision applications require the tracking of objects of interest. These objects of interest are sometimes stationary, which can complicate attempts to identify them. The goal of this project is to intelligently identify vehicles in video sequences under less than optimal conditions. While it is assumed that the video sequences are obtained from a stationary camera, algorithms robust to small camera movements are desirable. The algorithm used should also have low computational complexity. This paper focuses on both the recognition of objects of interest, as well as object correlation from one frame to the next.

Many methods, such as Background Subtraction and Corner Detection, have been proposed to perform this task. Background Subtraction identifies moving objects by comparing each video frame with a calculated background model. Alternately, Corner Detection finds the right angles in the images and tracks their movement to represent the moving objects. In this project, we propose an alternate variance-based algorithm. The speeds, locations, and intensity distributions of the objects are then used to assist in frame-to-frame correlation. The effectiveness of this variance-based algorithm is also compared to that of Background Subtraction and Corner Detection.

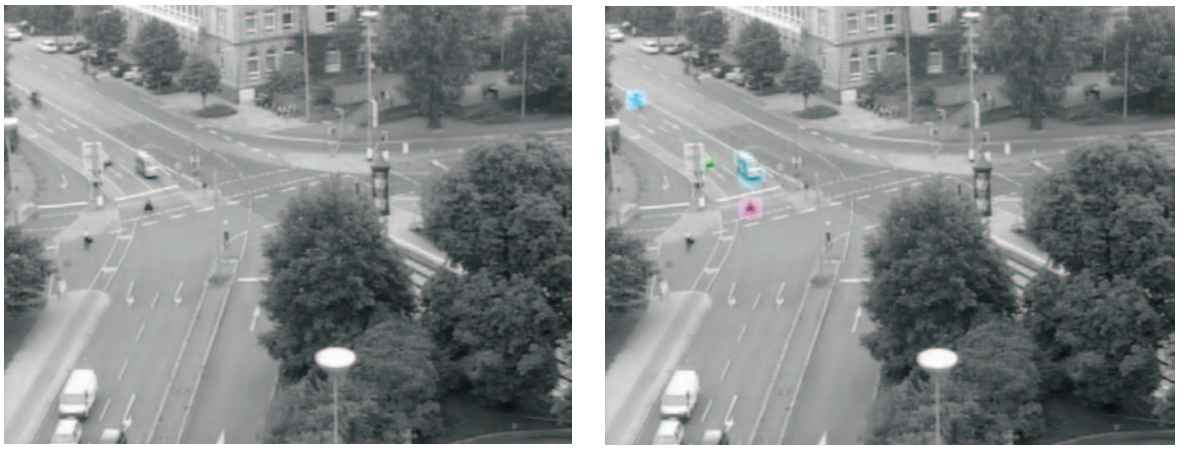

A scene from a video of a traffic sequence (left image), with the moving objects highlighted (right image). The moving objects were obtained using the variance-based method.

Multigrid Prolongation Based on Sharp Convergence Theory

Student **Rolf Ryham,** Pennsylvania State University

Mentor **Rob Falgout,** CASC

I was assigned to a project initiated by Visiting Scholar Prof. Ludmil Zikatanov and CASC faculty Panayot Vassilevski under the supervision of Rob Falgout. The project entailed the study, implementation, and presentation of an optimal prolongation algorithm for a two-level method based on a theory developed by the above three authors. An expose of the algorithm will be given below. Prolongation is one of three major user- or problem-defined components (prolongation, relaxation, coarsening) in a multigrid algorithm and is thus an integral part to the Linear Solvers Group's development of scalable solvers for large linear systems.

The nonconvexity of the algorithm makes it unclear whether certain key approximation steps are convergent or even stable, and we did not succeed in implementing a convergent algorithm for one- or two-dimensional problems. However, many modifications and test cases have been postulated for further development and will be implemented in the future. The prolongation operator P is mapping from the lower-dimensional coarse space to the high-dimensional fine space whose span determines the coarse basis. The coarse basis should reasonably represent the error after a smoothing operation, the essence of which may be expressed in the sharp bound:

$$\rho(E) = 1 - \frac{1}{K}, \quad K = \sup_{v \in R^n} \frac{\left\| (I - \Pi_{\tilde{M}})v \right\|_{\tilde{M}}}{\left\| (I - \Pi_A)v \right\|_A}$$

Choosing an optimal prolongation means to minimize K with respect to those P that satisfy a (linear) sparsity and an approximation constraint. The existence of a P is guaranteed but may not be unique. The algorithm postulates that an optimal P is a fixed point of the following iteration

1. Choose a high frequency v to represent the coarse grid complement.
2. Choose P that minimizes $\left\| (I - \Pi_{\tilde{M}})v \right\|_{\tilde{M}}$.

These steps are more difficult than the original matrix inversion. However, performing these approximately is feasible at the cost of knowledge of stability.

The AST Query Mechanism and the C/C++ Graphing Mechanism

Student

Andreas Sæbjørnsen, University of Oslo

Mentor

Dan Quinlan, CASC

I worked on two sub-projects of the ROSE project in cooperation with my supervisor Dan Quinlan. The sub-projects were the AST Query Mechanism and the C/C++ Graphing Mechanism, which are tools to extend and add capabilities to ROSE. They aim at increasing the productivity of the library developer using it.

The AST Query Mechanism is introduced as a building block when implementing algorithms and conditions in the ROSE abstract syntax tree (AST). When implementing a preprocessor, the user experiences the AST Query Mechanism as an interface to easily define and perform a subset of queries. This subset will typically be used to ask for a standard library list satisfying certain conditions, and it is possible to query for this in both a sub-tree of the AST and a list of SgNode pointers.

An example of a query is to ask for all Variable Declarations in the AST. The AST Query Mechanism introduces a method to easily ask the AST for Nodes satisfying conditions specified by the user, and also implements a set of predefined queries. The three different queries are the Name Query, Node Query and the Number Query. They are respectively returning a list of STL strings, a list of SgNode pointers and a list of integers. They all ask questions on all nodes in a sub-tree of the AST but are interested in different kinds of information. Because the AST Query Mechanism can query both a sub-tree of the AST and a list of SgNodes, it is often useful to combine the different queries to form an algorithm. For instance, one may ask the Node Query to create a list of all the Class Declarations in the AST, and another may ask the Name Query to query this list and return the names of all the classes.

The C/C++ Graphing Mechanism is an easy way to statically graph C/C++ data structures and provide a way to automatically generate code to dynamically graph the initialized data structures of a computer program at runtime. A graph of a computer program is a picture representing the relationship between objects and elements in the objects. For the user, the C/C++ graphing mechanism is a preprocessor that outputs a graph of the static data structure and the code needed at runtime to dynamically graph the code. When statically graphing the data structure, the data elements, such as variable declarations and class declaration, are graphed, and the static inheritance relationship is graphed. Dynamic graphing also graphs the data elements but is not graphing the inheritance relationship. The relationships that are graphed at runtime are the objects initialized to a pointer or reference and the values assigned to the defined variables. A dynamic graph is a snapshot of the memory allocated by the program at a specific time in the execution.

The AST Query Mechanism and the C/C++ Graphing Mechanism has successfully added value to the ROSE preprocessor project. Through these tools, the library developer can increase productivity while writing a preprocessor and debugging a program.

TOC

Implementation of OpenMP Support in the gFortran Compiler

Student **Elmer Salazar,** California State University, Stanislaus

Mentor **Mike Kumbera,** DCOM

OpenMP is a specification for a collection of compile-time directives and runtime functions in C and FORTRAN. This collection serves as a parallel programming interface. By using a standardized specification, the platform-specific details, as well as certain complexities of parallelism, can been hidden from the programmer. This allows for cleaner, more portable code, and compiling as parallel code is optional. If the programmer chooses to not compile as parallel code, the directives are treated as comments (ignored). This project's goal is to implement support for the OpenMP standard into the gFortran compiler. The gFortran compiler is a FORTRAN compiler that is part of the GNU GCC package. By implementing this standard into gFortran, programmers have access to OpenMP features, as well as the ability to see the intermediate code, including the transformations made by the compiler to handle parallelism.

Extending ReiserFS for Automatic File Queuing

Student

Jennifer Sirp, California State University, Sacramento

Mentor

Terry Brugger, NAIC

Typically, file systems are not concerned with directory entry organization. Tools like *ls* and *sort* are used to organize directory contents, but these tools perform linear operations that are costly when directories grow large. If files were maintained in the order they were written and stored in a priority queue, then a simple system call could return the "next file" to be processed. This summer, the benefit of adding automatic file queuing functionality to ReiserFS was explored.

A priority queue was implemented as an array of linked lists and maintained in kernel memory. As files are written on the partition, their names are examined to determine their priority and are then stored in the queue for rapid access. Inserts and retrievals are performed at a constant rate and priority queues exist for each directory that is created on the partition. Statistics on the queues are published through the proc File System, and access to the next file is achieved through a call to the *ioctl* function.

Tests were conducted on the new file system. A significant performance increase was achieved over prior user space applications that relied on multiple system calls and directory listings. Currently, work is being pursued on preserving the queue contents in hidden files so the data can be recovered in the event of system failure.

Efficient Schemes for Hyperbolic Systems with Stiff Relaxation Source Terms

Student **Yoshifumi Suzuki,** University of Michigan, Department of Aerospace Engineering

Mentor **Jeffrey Hittinger,** CASC

The objective of this study was to understand and develop numerical schemes that accurately and inexpensively solve hyperbolic systems with stiff relaxation source terms, systems that commonly arise in non-equilibrium hydrodynamic and radiation hydrodynamic problems. The generalized hyperbolic heat equations (GHHE) were used as a model system to investigate

$$\partial_t u + \partial_x v = 0$$
$$\partial_t v + \partial_x u = -\frac{1}{\varepsilon}(v - ru)$$

numerical schemes. The desired property of numerical methods for this problem is asymptotic preservation (AP), that is, the ability to accurately compute the diffusion limit when the relaxation process is underresolved. For this property, the scheme must account for the coupling between the advection and relaxation operators.

I briefly investigated the gas-kinetic scheme developed by Kun Xu. The idea of this BGK scheme is to evolve an approximation of the particle distribution function in each cell from which the macroscopic flow variables are updated. This scheme can accurately compute the Navier-Stokes limit by solving a Euler-like system provided that the distribution can support departures from equilibrium, but it is not clear how to extend this approach to more general systems.

Next, I investigated a particular semi-implicit Discontinuous Galerkin (DG) method [1] that has been shown to have the AP property in 1D. Accuracy in the diffusion limit is maintained because the solution slope is evolved directly with equations that account for the effect of the source term. I verified numerically that this scheme has the AP property in 1D by conducting convergence studies that compared to an exact solution. However, the AP property may be lost in higher dimensions. I have written a 2D DG code for a 2D extension of the GHHE system to investigate this issue, but I have not yet completed the convergence study. Investigating the 2D AP property of the DG scheme both numerically and analytically is the basis for future work.

Publications

[1] Lowrie, R.B. and J.E. Morel, "Methods for hyperbolic systems with stiff relaxation," *Int. J. Numer. Meth. Fluids* **40**, 2002, pp. 413-423.

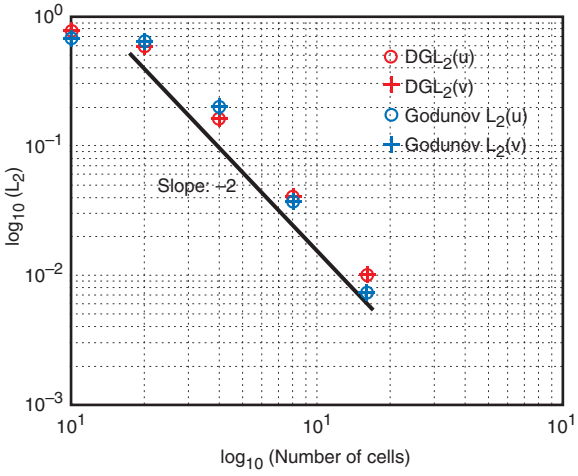

Figure 1. Grid convergence study of the DG method and the finite-volume Godunov method in the advection limit $(\varepsilon = 10^3, r = 0.5)$ solving the GHHE. This figure shows that both methods have second-order accuracy in this limit.

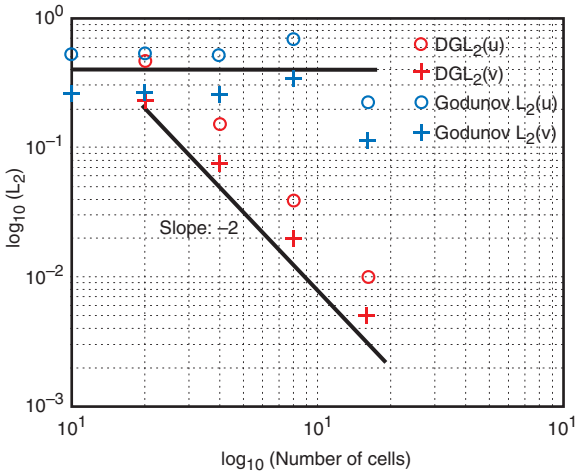

Figure 2. Grid convergence study of the DG method and the finite-volume Godunov method in the diffusion limit $(\varepsilon = 10^{-5}, r = 0.5)$ solving the GHHE. The DG method maintains second-order accuracy in this limit (AP property) unlike the Godunov method.

TOC

Streaming Computation of Structural Graphs

Student

Valerie Szudziejka, University of California, Davis

Mentor

Valerio Pascucci, CASC

Contour trees are useful for describing data simply while preserving extrema. We have implemented an efficient algorithm for constructing contour trees on two-dimensional data sets. This approach is easy to implement, and unlike previous algorithms, does not require us to sort any data, and scales to large data sets.

Our algorithm is expected to run in linear time in the number of critical points. One difference between our algorithm and previous algorithms is that we follow ascending edges of a mesh rather than following the isolines. It is simpler to label regions but more complicated to connect up the contour tree.

The colored figure below is a simple region where height corresponds to its y-coordinate. Maxima and minima are blue and saddles are black. Red points are not true extrema, they are 'starting points' that arise because our algorithm follows ascending edges of the mesh only.

Ordinarily, branches of a contour tree connect at a saddle point. In our method, branches connect along the line that connects the false extrema to its saddle. The final contour tree is equivalent—the stick picture on the right.

Future work includes extending this algorithm to three-dimensional data sets and adapting an out-of-core hierarchical representation of the data set such as the Z-order space-filling curve.

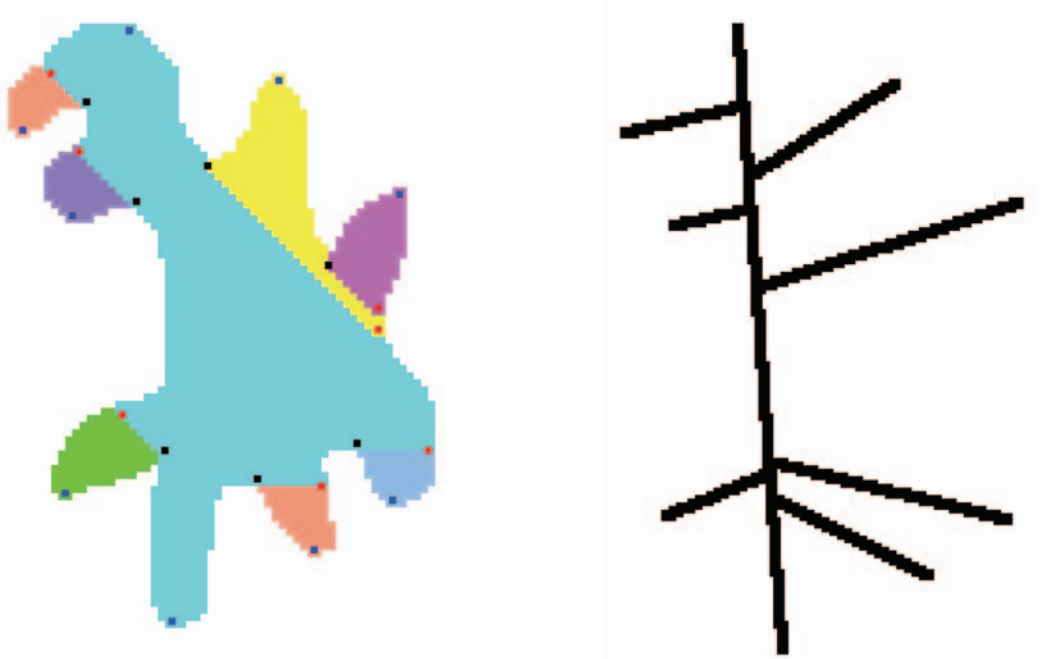

The colored image (left) is encoded with Z-order space-filling curves. The corresponding structural graph (right) is called a Contour Tree and is obtained by contracting each region of a distinct color to an arc in the tree. The computation is redesigned in a non-conventional way to achieve high performance in a streaming infrastructure.

TOC

Parallel AMG for Systems of PDEs

Student **Ryan Szypowski,** University of California, San Diego

Mentor **Ulrike Yang,** CASC

Algebraic multigrid (AMG) is a linear solver applicable to a wide variety of problems. However, when the linear system to be solved is derived from a system of coupled partial differential equations (PDEs), standard AMG often performs unsatisfactorily.

In such cases, it is necessary to augment the standard AMG scheme using knowledge of which unknowns are associated with a single node in the discretized PDE. One approach, known as the unknown approach, is to ignore connections between the different functions and coarsen them separately. Another approach, known as the nodal approach, is to categorize nodes of the discretization—as opposed to unknowns—as coarse or fine nodes, thereby keeping the coupled structure through all levels of the multigrid hierarchy.

The nodal approach has been studied in past work, and for small elasticity problems, it was seen to show an improvement over standard AMG. To test this method on larger problems, it was necessary to extend code previously written to work in the parallel version of Hypre. A large amount of code was added to Hypre to accommodate parallel block-structured sparse matrices. Also, special interpolation and relaxation routines were developed for use with the block-sparse matrices.

Although the code is not yet complete, we have hopes that this method will work well for a wide range of elasticity problems.

Computation of Cellular Detonation

Student

Brian Taylor, University of Illinois, Urbana-Champaign

Mentor

Bill Henshaw, CASC

A detonation is a region of rapid chemical reaction in which a very strong shock wave is sustained by the reaction heat release. Detonations can occur in explosives that are premixed gases, liquids, and solids. The nominal speed of a steady detonation wave is experimentally found to be close to the Chapman-Jouguet (CJ) detonation wave speed, which corresponds to a case where the state at the end of the reaction zone is sonic relative to the steady frame. Typically, the lead shock dynamics are influenced only by a finite region, mostly in the reaction zone.

Detonations are subject to strong multidimensional, cellular, and chaotic instabilities associated with shock-shock dynamics. Shock collisions may generate disturbances in the reaction zone. Typically, two portions of a detonation shock that collide will generate an oblique shock interaction. The oblique shock can generate reflected shocks and slip lines, which induce short-lived Kelvin-Helmholtz instabilities in the region near the detonation front. In two dimensions, the steady detonation is unstable and develops a cellular structure.

My work this summer involved simulating cellular detonation with a simple one-step reaction mechanism. I worked with a code developed at UIUC and with OverBlown, an Euler solver written using the Overture PDE framework. Primarily, I was interested in applying adaptive mesh refinement (AMR) to this problem using OverBlown, as this is a feature that we would like to incorporate into the code my research group at school is developing. I also spent quite a bit of time looking at how object-oriented programming can best be used in numerical programs. OverBlown/Overture is a large object-oriented C++ code. The design philosophy underlying it was quite instructive to learn about, especially the ways in which C++ is used to organize and encapsulate data with Fortran implementing the core numerics.

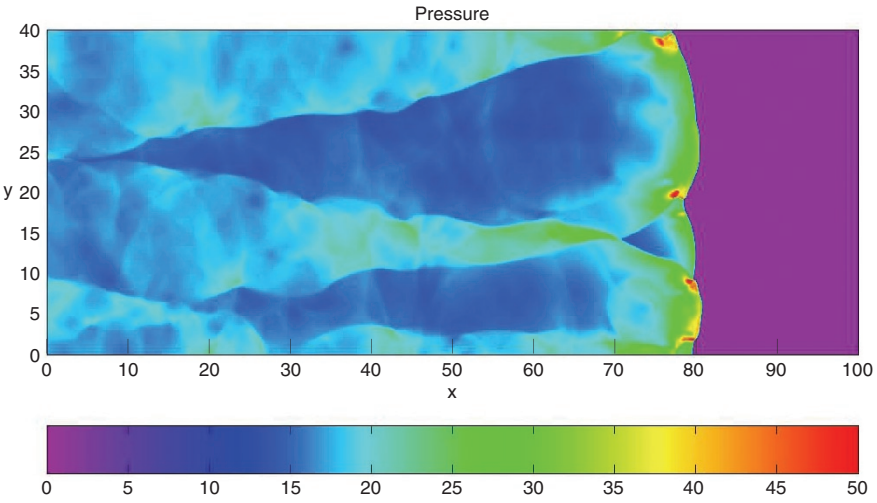

Cellular Detonation Pressure Profile

Vulnerability Tracking Database 2.0

Student

Rober Taylor, Northern Arizona University

Mentor

Jody Malik, CSP

The Vulnerability Tracking Database (VTDB) 2.0 project involved a rapid-development parallel redesign and upgrade of the current vulnerability tracking database and database interface scripts. In order to fully develop the new version of the database and interface within the short time period, I worked directly with Barry Dahling (CSP). Mr. Dahling was responsible for the major database modifications and redesign, while I was responsible for the interface and database interactions.

The current interface, implemented in single 3000-line Perl file, was redesigned and rewritten to create an easy-to-maintain PHP file hierarchy. For the interface, a new layout using Cascading Style Sheets (CSS) was created to allow for a simple, yet elegant, HTML interface to the new database. In addition to the redesign, many requested components were implemented in the interface. A knowledge base was added to the interface to provide vulnerability solutions and comments. The single load was upgraded to a bulk loader to increase productivity by reducing the number of repetitive tasks. Last, a new host report was created to provide vulnerability and comment history of a given host.

One of the main concerns with the current interface is access control. Therefore the VTDB 2.0 and accompanying interface were designed with access control in mind. Finer granularity was added to the VTDB, allowing for the designation of read and read-write users, along with finer directorate control. For the interface, I implemented an authorization component to control access to the interface and perform directorate control and checks.

Global Analysis of the ROSE Infrastructure

Student

Nils Thuerey, University of Erlangen-Nuerenberg, Germany

Mentor

Dan Quinlan, CASC and **Jody Malik,** CSP

The ROSE framework allows the transformation of C++ source code taking into account the semantics of higher-level abstractions. As its use is intended to be similar to a standard compiler, the need for global analysis results for a project arises. Hence, the first objective was to introduce a database that is able to store that information to ROSE. Furthermore, callgraph and classhierarchy traverals were chosen to test the database and simplify further program analysis within ROSE. As the callgraphs of typical programs turned out to be large, a tool for interactive visualization of graphs with many nodes and edges was written.

For the database management system, MySQL was chosen due to its powerful features and its availability as Open Source software. Additionally, the C++ interface to the database MySQL++ is used. The callgraph and classhierarchy generation programs were implemented as ROSE preprocessors. Both graphs use the BOOST graph library to simplify common tasks. While the classhierarchy generation was straightforward, callgraphs in C++ have to handle several language specific issues.

Function pointers and virtual functions make callgraph generation for C++ difficult. Virtual functions were handled by using the classhierarchy to identify virtual functions and possible calls to inheriting classes redefining these functions. For exchanging graphs with other programs, the .dot format from the GraphViz package is used within ROSE. However, as the Graphviz tools were not able to conveniently display larger graphs, a tool using the hardware acceleration of modern graphics cards to display and traverse graphs was written (dotgl). It uses the Graphviz development libraries to layout graphs and generate positions and sizes of nodes, edges and labels in the graph. The graph is then displayed with OpenGL. Due to graphical optimizations, such as adaptive level of detail and view frustrum culling, even graphs with several thousand nodes can be displayed in realtime.

The database interface, callgraph and classhierarchy are part of the ROSE framework now and will be used in future program analysis. The dotgl graph visualization will be maintained as a sourceforge project. Future work on the callgraph will handle function pointers, and evaluate its use for the automatic generation of components.

Publications

D. Quinlan & M. Schordan, "ROSE: A Tool for building Source-to-Source Preprocessors," 2003.

Y. Chiricota & F. Jourdan & G. Melancon, "Software Components Capture Using Graph Clustering," 2003.

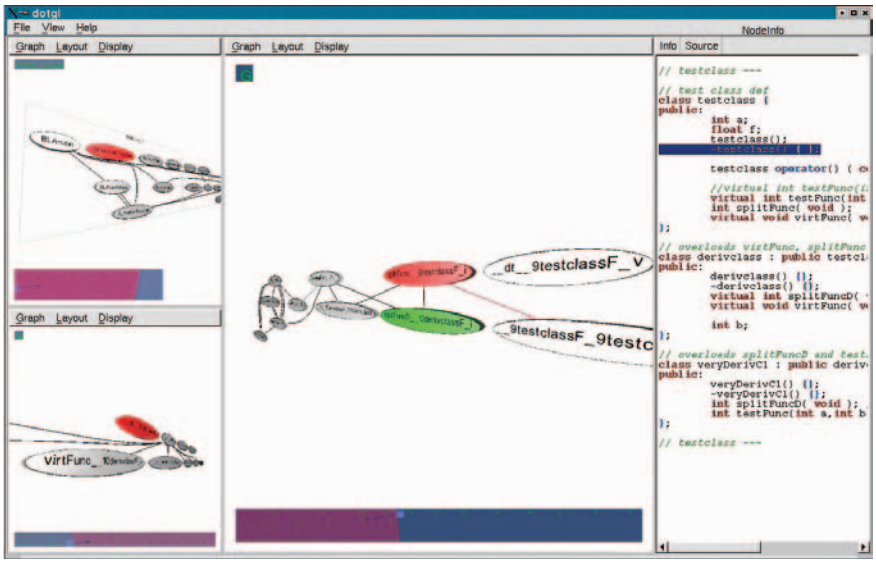

Three different callgraphs generated with ROSE and displayed with dotgl are shown. The sourcecode for the selected node of the center callgraph is shown in the panel to the right.

Code Validation Made Easy

Student **Peter Tipton,** University of Southern California

Mentor **Brandon Scott,** AX

Using DAKOTA (Design Analysis Kit for Optimization and Terascale Applications), users can perform a variety of studies to see how their code runs under different conditions. Among other things, users can run parameter studies, uncertainty quantification, and optimization. Using these tools, users can see how their code's output changes as their input parameters are varied, semi-automating the process of code validation.

The goal of this project is to take DAKOTA and be able to run it on an arbitrary code on an arbitrary machine. This should be done in such a manner that users will be able to analyze their code in the easiest, most powerful way possible. Our implementation of this technique is done with Python scripts.

A MATLAB Implementation of Mixed Finite Element Method for Incompressible Newtonian Flows: Pseudostress-Velocity Formulation

Student **Chunbo Wang,** Purdue University

Mentor **Charles Tong,** CASC

We investigate the use of a mixed finite-element method based on the pseudostress-velocity formulation for simulating incompressible Newtonian fluid flows. Semi-implicit time discretization of the formulation leads to a decoupled system at each time step such that one can calculate the pseudostress first and then the velocity. Physical quantities such as stress, vorticity and pressure can all be computed algebraically from pseudostress. The use of Raviart-Thomas finite elements in discretizing the decoupled pseudostress equation gives a linear system of equations that can be solved efficiently by a multigrid-preconditioned conjugate gradient method.

We developed a MATLAB code to show this scheme applied to the Lid-Driven Cavity problem. We also verify the accuracy of this scheme numerically by considering a 2-D Stokes problem. A GMRES preconditioned with H(div) geometric multigrid method is also developed.

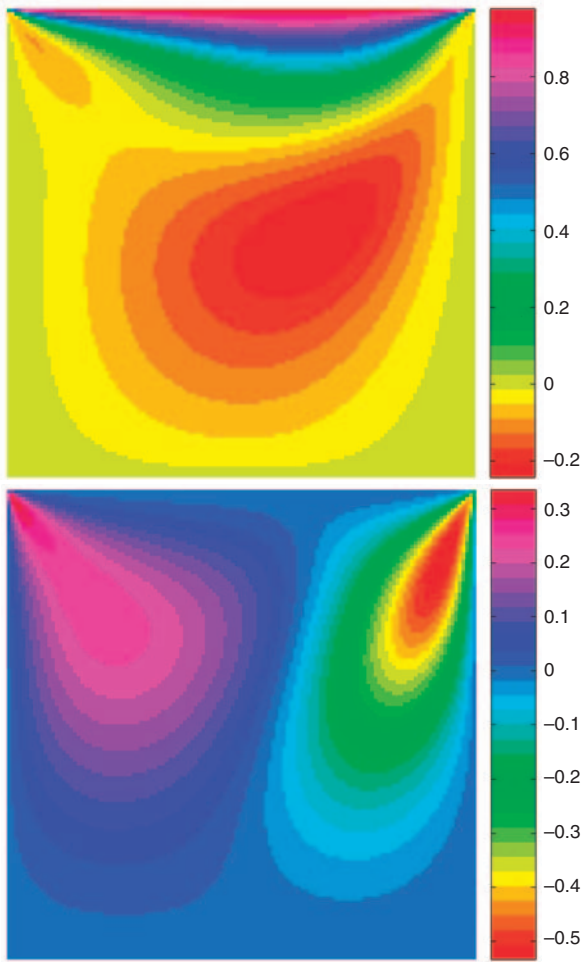

Plot of velocity of Lid-Driven Cavity problem

Newton–Krylov Methods for Expensive Nonlinear Function Evaluations

Student **Rebecca Wasyk,** Worcester Polytechnic Institute

Mentor **Carol Woodward,** CASC

Newton-Krylov methods have proven useful for solving large-scale nonlinear systems. An advantage of these iterative methods is that they do not require storage of the system Jacobian, but only require knowledge of how the Jacobian acts on a vector. A difference quotient evaluated at each linear iteration is often used to approximate this action without slowing the convergence rate of the method. For systems with expensive nonlinear function evaluations, however, the requirement of a function evaluation for each linear iteration can result in a very costly computation. The goal of this project is to explore convergence rates and time savings associated with using different approximations to the system function in the difference quotient.

Theoretical results developed in cooperation with Peter Brown, Homer Walker, and Carol Woodward indicate conditions on an approximation function that could be used instead of the full nonlinear function in the difference quotient without negatively impacting the local convergence rate. One such approximation that satisfies these conditions is approximating the most expensive parts of the nonlinear function with a linear approximation. Since making a linear approximation could be expensive, we also tested a method where the most expensive nonlinearities were lagged, so that they did not need to be recomputed in the difference quotient at each linear iteration.

These two approximations schemes were tested on several problems using the KINSOL solver in the SUNDIALS program, a suite of solvers developed at LLNL for solving nonlinear algebraic systems, ordinary differential equation systems, and differential-algebraic systems. For most of the problems tested, the linear approximation method did converge in the same number of iterations as were required when using the unmodified difference quotient, though in a few cases failures resulted when using this approximation. When the lagged approximation worked, we saw significant drops in the runtime in comparison to the unmodified method. As predicted by theory, the linear approximation had the same convergence properties as using no approximation and there were some noticeable time savings using this linear approximation on a few problems. However, none of the test problems had the type of expensive nonlinearities often seen in real application codes, where we would expect to see the most time savings.

Future work aims to develop some theoretical results on solving time-dependent problems using these approximations in the difference quotient. More tests are also planned on problems with more expensive nonlinearities, where the modifications made could really be beneficial in terms of runtime savings.

Metadata Management for Petabyte-Scale File Systems

Student **Sage Weil,** University of California, Santa Cruz

Mentor **Tyce McLarty,** CADSE

Previous simulation work validated many of the basic design points of a dynamic, distributed, subtree-based metadata server for large object-based file systems. Performance of the metadata server in file systems that decouple metadata from read and write operations is critical to overall system performance and scalability. Real-world performance of the design at scale remains speculative, due to the small scale of past simulation and the impracticality of simulating metadata operations at scale (an order of magnitude larger than existing Lustre installations).

To further study metadata performance issues, I have created a fully distributed prototype of the metadata server. The system is based on an asynchronous (but single-threaded) programming model based on C++ and STL, utilizing class-based continuations for delaying execution. Most of the basic functionality has been implemented, including a generic messaging layer (currently utilizing MPI), dynamic load balancing, synthetic client workload generation, and most major metadata operations needed to support the POSIX file system interface.

Information Leakage due to Geographic Properties of Internet Routing

Student **Daniel Wendlandt,** Stanford University

Mentor **Martin Casado,** ISCR

While it is commonly known that Internet routes operate with no restraint toward geopolitical boundaries, no empirical studies to date have explored how this notion of a network without borders impacts a nation-state's ability to limit access and control over its Internet traffic.

We sought to explore the issue of information leakage caused by geographic properties of Internet routing. We define information leakage to be when a route with both its source and destination in a single country uses infrastructure in another nation for transit. A better understanding to information leakage is crucial for national security concerns over controlling access to vital information and services running over the Internet. Our goal was to define a problem space, share data results, discuss underlying causes of information leaking paths, and explore diverse topics for future work in this area.

We devised a simple methodology for gathering introductory data on information leakage by using geolocation software to analyze traceroute paths between source/destination pairs within a single country. Our data results are compelling not only in that they are the first attempt to quantify the amorphous nature of information leakage but also because even our lower-bound results show that even for many highly advanced countries, the degree of information leakage is strikingly high. Additionally, our analysis shows that these circuitous paths are highly dependant on the type of service provider offering access to both the source and destination. We also explore how highly fiber-rich regions called "routing hubs" affect Internet routes.

TOC

Ameliorating the Performance Degradation of User-Defined Abstractions and Indirect Memory Accesses

Student **BrianWhite,** Cornell University

Mentor **Daniel Quinlan,** CASC

The complexity of physical simulations performed at LLNL demand high-level abstractions, such as vectors, tensors, and fields. While such abstractions facilitate program correctness, maintainability, and readability, they often inhibit performance because a compiler lacks sufficient semantic understanding to optimize them. For example, we encountered a field vector implementation based on a standard template library vector. Though the length of the field vector was known statically, it was not available to the compiler. Replacing the STL vector with scalars removed indirection inherent in the implementation and the overhead associated with its iterator and resulted in a 2x speedup.

Mesh-based codes exhibit indirect memory accesses, for example, as they visit nodes associated with an edge. Lab codes utilize iterators to mask such indirection. Unfortunately, the required indirection maps and iterator incrementing and dereferencing contribute to an excessive branch misprediction rate. One code suffered a 12.23% misprediction rate when measured on a Power3, achieving an anemic 0.88 instructions per cycle as a result on a machine with peak performance of 8 CPI (clocks per instruction). These unpredictable branches appear to be the performance bottleneck; a data-packing optimization that reduced L1 misses by 18% had no discernable effect on performance given the misprediction rate.

Despite the performance degradation induced by high-level abstractions, we believe they are necessary for effective programming. We seek to provide the performance benefit of more efficient, but less succinct, code through automatic transformation within the ROSE framework. Automating some of the above techniques will serve as the starting point for such future work.

Matching Shapes Using Local Descriptors

Student **Ryan White,** Univeristy of Calilfornia, Berkeley

Mentor **Shawn Newsam,** CASC

We present a method for comparing shapes of grayscale images in noisy circumstances. By establishing correspondences in a new image with a shape model, we can estimate a transformation between the new region and the model. Using a cost function for deviations from the model, we can rank resulting shape matches. We compare two separate feature detectors—Scale Saliency and Difference of Gaussians. We show that this method is successful in comparing images of fluid mixing under anisotropic geometric distortions and additive Gaussian noise and that Scale Saliency outperforms the Difference of Gaussians.

This project is further described at the Sapphire Web site:
http://www.llnl.gov/casc/sapphire/sapphirehome.html

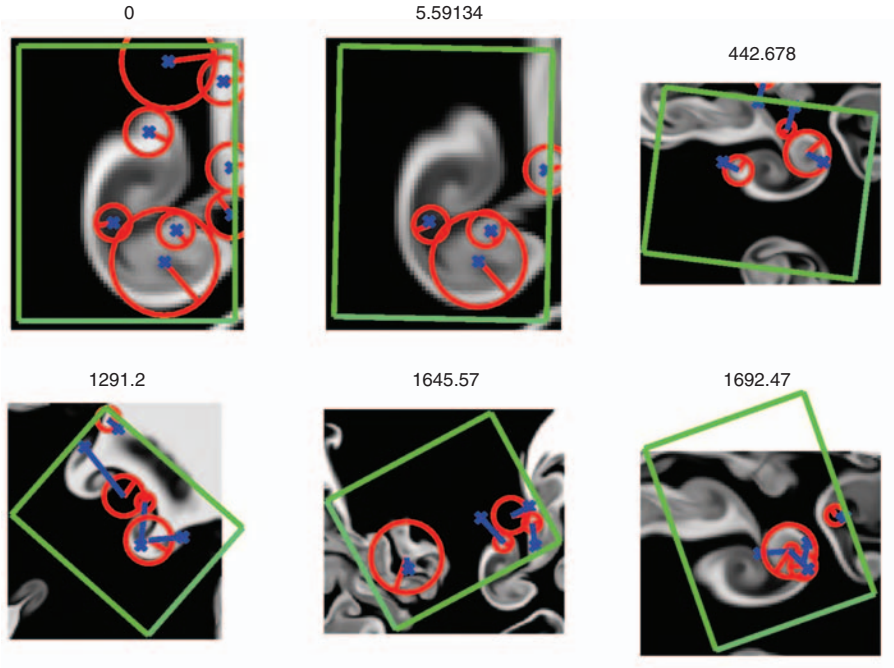

The green box in the top image represents the query region. The remaining images show the best matches in images that have been geometrically altered. The red circles are the features.

Additions to the ROSE Compiler Infrastructure

Student

Jeremiah Willcock, Indiana University

Mentor

Daniel Quinlan, CASC

My project this summer was to add to the ROSE project, a framework for creating source-to-source transformations for C++ programs. This project represents C++ programs as abstract syntax trees (ASTs) and provides features to parse C++ programs into ASTs, as well as to convert from ASTs back into C++ source code. I added several transformations on ASTs to the collection provided by ROSE.

One optimization I implemented was function inlining. This transformation replaces a call to a function with the body of the function, saving the overhead of the call, as well as enabling other optimizations on the combined source code. I also implemented two other optimizations using the ROSE framework—partial redundancy elimination and finite differencing. These standard compiler optimizations remove redundant computations and simplify complex computations. All three of these transformations are now included in the current version of ROSE, so ROSE users can also employ these transformations as part of their own optimizations.

Recently, I prototyped an interface between the Sage AST format used by ROSE and the Aterm file format for representing tree structures. This interface will enable the use of advanced rewriting tools, such as Stratego, within ROSE. In addition to my code additions to the ROSE project, I also contributed many bug reports and suggestions for improvement. Overall, I added several components to the ROSE project during my practicum this summer; these components will make ROSE more useful both for the Laboratory and for its external collaborators.

Shallow Water Equations on Curvilinear Grids

Student

Suzanne Wingenter, San Diego State University

Mentor

Petri Fast, CASC

This project solves the curvilinear version of the shallow water equations as described by Borthwick and Barber. [1] The equations were solved using an Adams Predictor Corrector scheme that was coded in C, compiled and run locally on the student workstation. This version was used to validate the finite-difference scheme using Richardson convergence rate studies. Convergence rates indicate the scheme accurately solves the equations with second-order convergence.

Three more test cases were prepared to check logical flow behavior. The first involves the development of flow in a straight channel over a small seamount showing the effects of bottom bathymetry on the surface elevation. Second, flow through a generic river geometry tests the ability of the flow in a contorted geometry. And third, flow past an idealized island tests recirculation. All three test cases show excellent behavior.

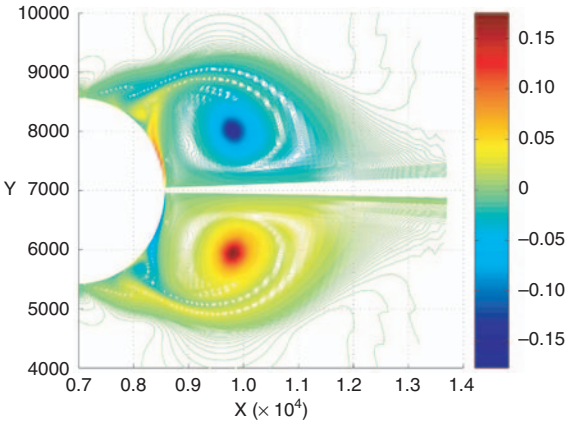

Flow past an idealized island shows recirculation, vorticity plot.

The last part of the project reimplements the code in Overture using native operators. Several simplified examples of the Bahia de Todos Santos (Baja California, Mexico) were run, the first of which is a smooth, coarse grid of the bay that can be run quickly. The other two examples were a flow past the northern most entrance to the bay, and a flow past the island of the bay. The last simulation performed involved the full bay. The simulation is running currently and the behavior is still in its transient phase.

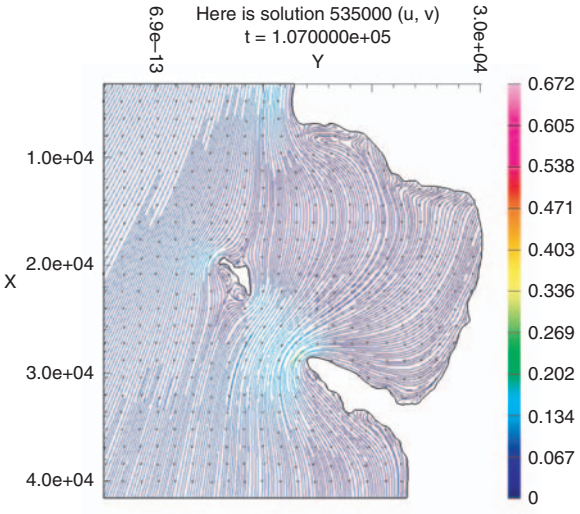

Transient behavior in Overture simulation of Bahia de Todos Santos, streamline plot showing beginning "steady state" tidal flow as the tide is starting to come in.

Publications

[1] A.G.L. Borthwick, RW. Barber, "River and Reservoir Flow Modeling using the Transformed Shallow Water Equations" *International Journal for Numerical Methods in Fluids*, 14:1193-1217, (1992).

A Hybrid Sort-First/Sort-Last Approach for Rendering Translucent Geometry in the VisIt Visualization Tool

Student

Christopher Wojtan, Georgia Institute of Technology (Fall 2004), University of Illinois, Urbana-Champaign (Fall 2000 – Spring 2004)

Mentor

Jeremy Meredith, DNT

The scientific visualization tool VisIt lacked any support for reading image formats as data sets. I implemented support for image file formats (PPM, BMP, JPEG, TIFF, etc.) within VisIt. The red, green, and blue channels within these images were treated as vector data and the pixel location was treated as the spatial location of the data. Then I created a way to plot the image in red, green, and blue colors, instead of using a variable color ramp to display each channel. VisIt did not handle translucency in two-dimensional data very efficiently, so I improved that, as well. The original scheme involved sorting all cells front to back, then calculating the transparency. Sorting was not necessary in two dimensions.

After this work, I began correcting the way VisIt displays translucent data when rendering in parallel. The original method assigned a group of data to each processor according to the data's location in world space. This led to problems with transparent data, since it needs to be sorted from front to back in image space. We decided to assign data to processors according to locations in image space in order to correctly perform this scalable rendering. Since the data was originally assigned incorrectly for our purposes, the processors needed to redistribute the data. I wrote code for each processor to transform its data into image space and redistribute everything correctly in parallel. Now the data is correctly rendered in each region of the screen, and we only need to consider clipping everything to finalize this scheme.

Cache Coherent Mesh Layout

Student **Sung-Eui Yoon,** University of North Carolina at Chapel Hill

Mentor **Peter Lindstrom,** CASC

Recently, streaming meshes have been introduced to efficiently perform out-of-core processing of gigantic meshes. The mesh is constructed such that memory requirements are minimized during processing the mesh following a streaming (or stored) order. Moreover, seamless connectivity of currently processing elements is provided. This streaming mesh greatly improves out-of-core performance of applications that do not require their own specific processing order. However, many applications require their own specific processing order, which is different from a stored streaming order.

To improve out-of-core performance of those applications, we need to have a mesh layout, which is stored in a specific order and, more importantly, maintain locality of mesh elements for cache coherent access pattern.

There is plenty of literature available on layouts of structured meshes (e.g., uniform grid). Classic space-filling curves have been used to compute ordering of mesh elements to maintain locality. However, embedding a space-filling curve into a uniform geometric structure can deteriorate quality of locality in unstructured meshes. To ameliorate this problem, we figured out an algorithm computing the ordering based on construction of classic space curves on unstructured meshes. Moreover, to quantify locality of the ordering, we analytically compute a probability of a cache miss according to edge length given cache parameters.

In the future, we plan to design a cache-oblivious algorithm that does not require exact cache parameters. This work will be carried out at University of North Carolina with Peter Lindstrom.

TOC

A Visual Interface for the Promoter Identification Workflow

Student **Beth Yost,** Virginia Polytechnic Institute and State University

Mentor **Terence Critchlow,** CASC

The next step in analyzing the human genome involves understanding gene regulation. Genes are turned on or off based on a complex series of mechanisms. These controls involve the promoter region of a gene, typically located right before the start of transcription. In this region, transcription factor binding sites can be found. Biologists determine promoter modules based on the frequency of these sites and the distance between them. By comparing these factors across multiple sequences, it is possible to determine other genes that might have a similar function.

Identifying these promoter regions is a difficult task. The Promoter Identification Workflow (PIW) automates this process, but a visual interface designed specifically for the biologist was lacking. We have begun developing an interactive visual interface for the PIW. This interface provides biologists with a quick overview of the results. Included in the overview is the ClustalW alignment of the sequences, shaded regions of those sequences representing the locations where there is homology across sequences, and colored boxes showing the locations of transcription factor binding sites (determined by TransFAC). A consensus sequence is calculated and is always shown at the top of the overview. There is a sliding window that shows details of 50 base pairs at a time. There are also filters based on the names and frequencies of transcription factor binding sites. Using the PIW with the added visual interface can not only speed the work of the biologist, but also provide them with additional insight.

ISCR Seminar Series

The ISCR hosted 70 seminars from visitors in FY 2004 covering a wide spectrum of research areas, and recruited an additional 35 speakers from LLNL ranks to speak to visiting students. The ASC Institute for Terascale Simulation Lecture Series was established in 2000 to enrich the intellectual atmosphere of LLNL's large simulation community through the visits of leaders representing the diverse areas of computation. In FY 2004, we hosted William Wulf, John Grosh, Bjarne Stroustrup, Mary Wheeler, and David Bailey. The ISCR Summer Student Lecture Series was also established in 2000 and forked into three different series in Summer 2004—Internships in Computational Modeling at the Terascale (ICMT), Internships in Computer Science at the Terascale (ICST), and College Cyber Defenders Computer Security.

TOC

December 4, 2003

Challenges for Computing and Information Technology in the 21st Century

Speaker

William. A. Wulf, National Academy of Engineering, wwulf@nae.edu

Information Technology (IT), the convergence of computing and communications technologies, has had an enormous impact on all aspects of life in the developed world. It will have even more impact in both the developed and developing world as we enter the 21st century. Powered by the unprecedented and continuing advances in microelectronics and photonics, the power and capacity of our expanding information infrastructure has risen exponentially while simultaneously its cost has also fallen exponentially. At least for the foreseeable future, the exponential pace of technology improvement is likely to continue.

In this lecture, I will explore some of the non-technical, societal challenges and opportunities posed by information technology as we enter the 21st century. I will not provide answers for these challenges or guarantees that we will exploit the opportunities, but hopefully, asking some of the right questions will provoke serious thought about them.

Institution Web page:
http://www.nae.edu/

TOC

February 17, 2004

A DoD Perspective on High-End Computing

Speaker

John Grosh, Department of Defense, john.grosh@osd.mil

Over the past several years, we have witnessed a renewed interest in high-end computing. The advent of the Earth Simulator and maturation of cluster technology have spurred intense discussion at both technical and high levels of government, industry, and academia on the future of high-end computing. Frequently, such discussions suffer from oversimplification, masking the real reason why we develop and deploy high-end computing, i.e., to advance science and engineering and to achieve specific operational capability.

In this talk, I will take a Department of Defense (DoD) perspective on this field, delving into the driving factors that will define the future of high-end computing. Central to this discussion are concepts of systems engineering, which are influencing DoD strategy in this area. I will also outline DoD activities, as well as technology and policy issues related to high-end computing. In particular, I will discuss performance measurement, importance of matching applications to systems requirements, software engineering, and a new area of concern, applications software security (as opposed to network and computer security). In addition, I will discuss some of the challenges of embedded computing systems and draw some loose parallels to high-end computing. In wrapping up the discussion, I will discuss what I view are the fundamental challenges facing high-end computing from both a technical and policy perspective.

TOC

March 1, 2004

Programming, Language and Libraries

Speaker

Bjarne Stroustrup, Texas A&M University, bs@cs.tamu.edu

We say we write our code in a programming language. For real code, that's only partially true. We write our programs in an "extended language" characterized by a set of libraries supporting our general application domains and key abstractions (e.g., graphics, linear algebra, and distribution). This talk explores the relationships between language features (e.g., classes, templates, and exceptions), programming styles (e.g., object-oriented programming and generic programming), and library design in C++. All examples will be very simple to illustrate fundamental programming and performance issues. The examples will illustrate some of the design principles of C++, and based on that, I'll make a few conjectures about the likely directions for the evolution of C++ and the International Organization for Standardization (ISO) C++ standards effort.

Speaker's Web page:
http://www.research.att.com/~bs/homepage.html

Institution Web page:
http://www.tamu.edu/

TOC

June, 24, 2004

Mathematical and Computational Modeling of Multiphysics Couplings

Speaker

Mary Fanett Wheeler, University of Texas, Austin, mfw@ices.utexas.edu

Multiphysics couplings can happen in different ways. One may have different physical processes (e.g., flow, transport, reactions) occurring within the same physical domain, or one may have different physical regimes (e.g., surface/subsurface environments, fluid/structure interactions) interacting through interfaces. We will discuss both of these types of multiphysics couplings during this presentation. Of particular interest will be the development of interpolation/projection algorithms for projecting physical quantities from one space/time grid to another, the investigation of mortar and mortar-free methods for coupling multiple physical domains, and the coupling of non-conforming and conforming finite element methods.

Speaker's Web page:
http://www.ticam.utexas.edu/~mfw/

Institution Web page:
http://www.utexas.edu/

August 4, 2004

Twelve Ways to Fool the Masses: Back to the Future

Speaker

David H. Bailey, Lawrence Berkeley National Laboratory, dhbailey@lbl.gov

In the early 1990s when highly parallel computing technology was new, researchers, as well as computer vendors, were prone to make inflated claims of the effectiveness and performance of these systems for scientific computations. Concerned by these developments, I published a tongue-in-cheek article in 1993 called "Twelve Ways to Fool the Masses When Giving Performance Results on Parallel Computers." Such warnings were not widely heeded, and as a result, the parallel computing field lost some measure of credibility, numerous parallel computer firms failed, and government agencies trimmed funding. Now the field is recovering, and new variants such as grid computing are emerging. But once again, researchers in the field are advised not to oversell this technology and to employ rigorous methodologies when analyzing and reporting results on these systems.

Speaker's Web page:
http://crd.lbl.gov/~dhbailey/

Institution Web page:
http://www.lbl.gov/

TOC

October 1, 2003

A C++ Reflection Library

Speaker

Kenneth Chiu, Indiana University, chiuk@indiana.edu

The ability to programmatically inspect a class definition can add power and flexibility to certain software systems. For example, a middleware library can use this to convert the state of an object into a sequence of bytes that can be transmitted over the network. This ability, termed reflection, relies on classes that contain metadata describing other classes. Some languages, such as Java, provide these classes as part of the language specification. However, C++ does not. In this talk, I present a C++ reflection library. I will cover both the public interface and some of the implementation specifics. I also survey some of the other existing work in this area.

Speaker's Web page:
http://www.cs.indiana.edu/~chiuk

Institution Web page:
http://www.iub.edu/

TOC

October 3, 2003

A User-Friendly Probability Machine with Applications

Speaker

Branden Fitelson, University of California, Berkeley, branden@fitelson.org

A general mechanical procedure with a user-friendly front end for reasoning about the probability calculus is presented. The procedure is then used to solve various probabilities in probability theory. Issues of computational complexity and problem size will also be discussed. All necessary technical and historical background will be provided during the talk.

Speaker's Web page:
http://fitelson.org

Research Web page:
http://fitelson.org/research.htm

Institution Web page:
http://www.berkeley.edu

TOC

October 9, 2003

Dynamic Code Generation in C++ as a Foundation for Domain-Specific Optimization

Speaker

Olav Beckmann, Imperial College, London, ob3@doc.ic.ac.uk

The TaskGraph Library is a C++ library for dynamic code generation that combines specialization with dependence analysis and restructuring optimization. A TaskGraph represents a fragment of code that is constructed and manipulated at runtime, then compiled, dynamically linked and executed. The TaskGraph Library is implemented purely in C++ using macros and operator overloading to define a simplified, C-like sub-language that is used for initializing TaskGraphs. The internal representation used for representing generated code is SUIF-1, and the TaskGraph library implements an API for calling SUIF's analysis and restructuring passes on the generated code.

We view the TaskGraph Library as a research tool for facilitating domain-specific runtime optimizations in scientific applications. One key distinction of this approach is its combination of runtime code specialization and restructuring optimization. Sample applications include:

- Specialization of a generic image-filtering application to a specific convolution matrix
- Automatic search for optimal tile size and loop sequence for matrix multiply and dense stencil loops
- Unrolling of loops over Morton-order matrices.

We are currently planning to use the TaskGraph library as a tool for implementing and evaluating domain-specific optimization "components" for scientific programs. We are particularly interested in exploring this idea in the context of stencil loops.

Speaker's Web page:
http://hpc.doc.ic.ac.uk/~ob3/

Institution Web page:
http://www.ic.ac.uk/

TOC

October 23, 2003

Progress in Clustering

Speaker

Charles Elkan, University of California, San Diego, elkan@cs.ucsd.edu

Clustering algorithms are widely used in data mining, but basic advances are still needed in order to build genuinely robust clustering software. In this talk, I shall describe three advances that together make the well-known *k*-means algorithm far more useful. The first advance is an optimization of *k*-means that uses the triangle inequality to avoid unnecessary distance calculations. Empirically, the optimized running time has almost no dependence on *k* and the number of clusters, so hundreds of clusters can be found in gigabyte data sets in minutes. The second advance is an extension of k-means that is far better at identifying true clusterings because it does not get stuck at local minima. Finally, the third advance is an extension of *k*-means that automatically finds the appropriate number of clusters to use. This algorithm uses a novel statistical test for the hypothesis that a data subset is Gaussian, which works regardless of the dimensionality of the data.

Speaker's Web page:
http://www.cs.ucsd.edu/users/elkan

Institution Web page:
http://www.ucsd.edu/

TOC

October 23, 2003

Statistical Models for Simulation Errors and Their Role in Prediction and Uncertainty Quantification

Speaker

James Glimm, State University of New York, Stony Brook, glimm@ams.suny.edu

Uncertainty in numerical simulations pertains to both forward simulation and to inverse problems. Probability estimates for simulation error are used to assess mismatch between approximate solutions and observations. The straightforward application of these ideas runs into a fundamental obstacle, namely an explosion of computational requirements. Determining whether an error needs a (better) comparison solution uses more resources, and determining a statistical ensemble of errors requires even more resources. To understand the parametric dependence of the statistics on the solution parameters also requires additional resources.

We present some ideas that may allow a simplification of the program, and thus allow its usefulness in practice:

1. Even deeply nonlinear problems can utilize linear models for the error statistics
2. Composition laws allow the build-up of error models for complex problems from simpler components
3. Parametric dependence of statistics can be expressed in simple regression models
4. Separate effects leading to uncertainty can be isolated and quantified individually. Methods to accomplish this step will be explained with illustration to shock physics and porous media simulations.

Our general message is that through scienfitic understanding of the errors, we will greatly reduce the required simulation resources needed for their characterization.

Speaker's Web page:
http://www.ams.sunysb.edu/~glimm/glimm.html

Research Web page:
http://www.ams.sunysb.edu/~shock/FTdoc/FTmain.html

Institution Web page:
http://www.sunysb.edu/

TOC

October 24, 2003

Simplification, Adaptivity and Conservation in Front-Tracking Method

Speaker

Xiaolin Li, State University of New York at Stony Brook, linli@ams.sunysb.edu

I will discuss three important issues in the front tracking method. First is the simplification of the geometrical handling—we use the grid-based interface reconstruction method to reduce the topological complexity. Second, we combine the tracking and adaptive methods. We developed the interoperability between two codes, FronTier and Overture. The former handles tracking while the latter handles adaptivity. We also address the conservation issue of the front tracking method through the use of the dynamical flux.

In addition, I will report on two recent developments of the front tracking method. One is the locally grid-based tracking method. This method combines the merits of both grid-free tracking and grid-based tracking. Another advance is the extension of the three-dimensional front tracking for multi-component interfaces.

Finally, I will show some success stories of the front tracking method when it is applied to important physics and engineering problems, such as the study of the Rayleigh–Taylor instability.

Speaker's Web page:
http://www.ams.sunysb.edu/~linli

Research Web page:
http://www.ams.sunysb.edu/~shock/FTdoc/FTmain.html

Institution Web page:
http://www.sunysb.edu/

TOC

October 27, 2003

Image Processing: From Segmentation to Recognition

Speaker

Achi Brandt, Weizmann Institute of Science, achi@wisdom.weizmann.ac.il

Image segmentation is a prerequisite for higher-level processes from motion detection to object recognition and can also be used for picture denoising, compactification, miniaturization, etc. Segmentation is difficult because objects may differ from their background by any of a variety of properties that can be observed in some, but often not all scales. Our algorithm of Segmentation by Weighted Aggregation (SWA) consists of an adaptive process in which pixels are recursively aggregated into increasingly larger scale aggregates, each having progressively longer lists of coherent properties, including statistics from sub-aggregates at all finer levels. The process, motivated by the algebraic multigrid (AMG), is enhanced by top-down (coarse-to-fine) feedback of aggregation directives.

Since the computation is recursive and done mostly at coarser levels, the algorithm costs only several dozen operations per pixel and is highly parallelizable. Experimental results demonstrate a dramatically improved segmentation over current state-of-the-art methods. Moreover, the hierarchical segmentation produced in this way is in a form directly usable by learning/recognition systems since each segment emerges equipped with a vector of numbers representing textures, standardized shapes, sub-segments with their own vectors of numbers, and other identifying features. (Joint work with E. Sharon, M. Galun, Y. Gorelick and R. Basri)

Speaker's Web page:
http://www.wisdom.weizmann.ac.il/~achi/

Institution Web page:
http://www.weizmann.ac.il/

TOC

October 30, 2003

Efficient Solution of the Discretized S_N Time-Dependent Boltzmann Transport Equation on Parallel Platforms

Speaker

Douglas Swesty, State University of New York at Stony Brook,
dswesty@blackhole.ess.sunysb.edu

Implicit S_N discretizations of the time-dependent Boltzmann transport equation give rise to large, sparse, linear systems of equations that require the use of parallel architectures. Two of the most common methods for solving these sets of equations are source iteration (commonly used in nuclear engineering) and a full-matrix approach (commonly used in astrophysics). I will present results of a comparison study of the efficiency of these two approaches on a variety of 1d and 2d test problems.

Speaker's Web page:
http://www.astro.sunysb.edu/dswesty/

Institution Web page:
http://www.sunysb.edu/

TOC

October 31, 2003

Detailed Cache Coherence Characterization for OpenMP Benchmarks

Speaker

Frank Mueller, North Carolina State University, mueller@cs.ncsu.edu

Past work on studying cache coherence in shared-memory symmetric multiprocessors (SMPs) concentrated on processor simulation, including the memory hierarchy and interconnects. With the advent of hardware counters, it has become feasible to study coherence traffic based on aggregates of events. But little work has been put forward to combine the benefits of simulation methods and hardware counters. Our work closes this gap with an approach for fine-grained coherence simulation indicating opportunities for optimizations that can subsequently be confirmed by evaluating hardware counters.

The technical contributions of this work are as follows. We introduce ccSIM, a cache-coherent memory simulator fed by data traces obtained through on-the-fly dynamic binary rewriting of OpenMP benchmarks executing on a Power3 SMP node. We explore the degrees of freedom in interleaving data traces from the different processors compared to hardware performance counters. The novelty of ccSIM lies in its ability to relate coherence traffic—specifically invalidations resulting in subsequent cache misses—to data structures and to their reference locations in the source program, thereby facilitating the detection of inefficiencies.

Our experiments demonstrate that:

(a) Cache coherence traffic is simulated accurately for SPMD programming styles as its invalidation traffic closely matches the corresponding hardware performance counters

(b) We derive detailed coherence information indicating the location of invalidations in the application code

(c) We derive opportunities for optimizations from these details leading us to program transformations that result in decreased coherence misses and execution time.

Speaker's Web page:
http://moss.csc.ncsu.edu/~mueller/

Institution Web page:
http://www.ncsu.edu/

November 7, 2003

Computational Models of Cellular Metabolic Fluxes

Speaker

Daniel Segre, Harvard Medical School, dsegre@genetics.med.harvard.edu

Evolutionary adaptation can result in biological systems that are optimal for certain tasks, compatibly with their physico-chemical constraints. In the absence of a detailed knowledge of all parameters for a metabolic network, optimality can serve as a powerful constraint to produce quantitative testable predictions. Based on the premise that bacteria such as Escherichia coli have maximized their growth rate during evolution, flux balance analysis (FBA) predicts whole-cell metabolic reaction rates (fluxes) at steady state using linear programming. However, while the assumption of optimal growth rate for a wild-type bacterium is justifiable, the same argument may not be valid for artificially perturbed strains, such as genetically engineered knockouts.

I will show how one can identify such suboptimal metabolic states by assuming that the immediate response to the perturbation is a minimization of metabolic adjustment (MOMA) with respect to the wild-type flux distribution. The study of flux differences between suboptimal and optimal metabolic network states can be useful in understanding regulatory and evolutionary responses to environmental changes and ecosystem dynamics.

Speaker's Web page:
http://genepath.med.harvard.edu/~dsegre/

Institution Web page:
http://hms.harvard.edu/

TOC

November 13, 2003

Linear Algebra and Quantum Simulations

Speaker

Richard Scalettar, University of California, Davis, rts@sherlock.ucdavis.edu

One of the most powerful methods for understanding the physics of many interacting electrons (magnetism, superconductivity, metal-insulator phase transitions, etc.) is the determinant quantum Monte Carlo algorithm. In this approach, the electronic degrees of freedom are integrated out through the introduction of an auxiliary field, leaving an expression for the partition function that consists of a very high-dimensional integral over the auxiliary field. The integrand is a product of determinants whose values depend on the field configuration. To do the integral stochastically involves computing how the determinant changes as the field configuration changes, which ends up requiring the repeated calculation of the inverse of the matrix whose determinant is the integrand.

In this talk, I will briefly review the determinant quantum Monte Carlo algorithm and then focus on the detailed structure of the matrices that arise. I will then discuss open questions of how to do the linear algebra involved more efficiently. In particular, we are looking at an alternate approach which in principle reduces the computation of the matrix M^{-1} to that of M^{-1} on a vector and thereby the scaling of the algorithm from the cube of the matrix dimension to linear. Unfortunately, the computation of M^{-1} on a vector goes horribly awry (no convergence) in practice.

The goal of this talk is to get input into the solution of these problems.

Speaker's Web page:
http://leopard.physics.ucdavis.edu/rts/

Institution Web page:
http://www.ucdavis.edu/

TOC

November 14, 2003

Cartesian Grids with Embedded Geometry

Speaker

Marsha Berger, New York University, berger@CIMS.nyu.edu

Cartesian mesh methods are one approach to solving partial differential equations (PDEs) in complex geometry. In this approach, only cells cut by the geometry require special attention. We review Cartesian mesh methods and discuss the problem of constructing stable and accurate discretizations at the cut-cells. Computational results for realistic three-dimensional aircraft will be presented.

This is joint work with Michael Aftosmis and Scott Murman of NASA Ames Research Center, Randy LeVeque at the University of Washington, and Christiane Helzel at the University of Bonn.

Speaker's Web page:
http://www.cs.nyu.edu/cs/faculty/berger/

Institution Web page:
http://www.nyu.edu/

TOC

November 24, 2003

Techniques for Processing Large Data Streams

Speaker

Johannes Gehrke, Cornell University, johannes@cs.cornell.edu

Data management techniques for data streams have gained much importance recently. I will talk about two techniques. First, I will introduce techniques for approximately answering queries over continuous data streams with limited memory. Our method relies on randomizing techniques that compute small "sketches" of a stream that can be used to provide approximate query answers with provable error guarantees. The second part of the talk will introduce techniques for efficient integration and aggregation of historical information for archiving data.

Speaker's Web page:
http://www.cs.cornell.edu/johannes/

Research Web page:
http://www.cs.cornell.edu/database/

Institution Web page:
http://www.cornell.edu/

TOC

December 3, 2003

Cray Cascade Project

Speaker

John Feo, Cray, Inc., feo@cray.com

In 2002, the US Department of Defense initiated the High Productivity Computing Systems Project to develop a next-generation computer system capable of sustaining a petaflop. Cray Incorporated was selected as one of five commercial vendors to submit an initial design. Cray, working with its partners, has proposed Cascade, a revolutionary new computer system comprised of custom processors, a next-generation interconnection network, and an active memory system.

While still in design, the system is expected to include support for both heavyweight threads that exploit high temporal locality and lightweight threads that exploit high-spatial locality. The former will execute on processors that tolerate memory latencies through a combination of multithreading, vector, and stream processing. The latter may execute in active memory systems with PIM-like characteristics that may also be multithreaded to tolerate memory latencies. The interconnection network may be a symmetric Cayley graph network capable of high bandwidth, low latency communications. Memory will be physically distributed, but shared.

A sophisticated programming environment is proposed to assist application programmers with automatically utilizing the machine's unique processing capabilities. We expect that the global shared memory and the hardware's ability to tolerate memory latencies when executing either heavyweight or lightweight threads will eliminate many of the programming challenges confronting scientific application developers today.

In this talk, I will present the design goals for Cascade and describe the architecture and programming environment as they are currently envisioned.

Institution Web page:
http://www.cray.com/

TOC

December 3, 2003

The Magneto-Hydrodynamic Richtmyer–Meshkov Instability

Speaker

Ravi Samtaney, Princeton Plasma Physics Laboratory, samtaney@pppl.gov

In the past two decades the Richtmyer–Meshkov (RM) instability has become the subject of extensive experimental, theoretical and computational research due to its importance in technological applications such as inertial confinement fusion, as well as astrophysical phenomena such as shock interactions with interstellar clouds. In this talk, we will present recent results from nonlinear simulations of the Richtmyer–Meshkov instability in the presence of a magnetic field.

The seminar will be divided into three segments. In the first segment, we will present a primer on compressible magneto-hydrodynamics (MHD). In the second segment, we will present numerical evidence that the growth of the Richtmyer–Meshkov instability is suppressed in the presence of a magnetic field. This is due to a bifurcation that occurs during the refraction of the incident shock on the density interface. The result is that baroclinically generated vorticity is transported away from the interface to a pair of slow or intermediate magnetosonic shocks. Consequently, the density interface is devoid of vorticity and its growth and associated mixing is completely suppressed. The third segment on the talk will focus on the numerical method to obtain the aforementioned results. We will discuss the implementation of an unsplit upwinding method to solve the ideal MHD equations with adaptive mesh refinement (AMR) using the CHOMBO framework. The solenoidal property of the magnetic field is enforced using a projection method that is solved using a multigrid technique.

Research Web page:
http://w3.pppl.gov/APDEC-CEMM

Institution Web page:
http://www.pppl.gov/

TOC

December 9, 2003

Why Are Streamline Methods Attractive for Simulation of Gas-Injection Processes?

Speaker

Margot Gerritsen, Stanford University, margot.gerritsen@stanford.edu

In the SUPRI-C research group, we are interested in the design of efficient and accurate simulation tools for compositional problems, such as those occurring in gas-injection processes. The underlying method we use is the streamline method. In this talk, I would like to motivate this choice and discuss recent extensions and improvements made to the traditional streamline method that improve its accuracy and efficiency for compositional problems.

For reliable performance prediction of gas-injection processes, it is essential to properly account for the effects of heterogeneity through the use of fine grids and accurately represent component transfer between phases. Because of the required high-grid density and the costly phase behavior calculations (i.e., flashes), gas-injection processes are computationally intensive. Streamline methods akin to Euler-Lagrangian methods are attractive for these advection-dominated processes because they alleviate time-step restrictions, are inherently parallel, and easily combined with adaptive mesh refinement strategies to reduce flashing costs. In streamline methods, the pressure equation is solved on a 3D grid, while the mass balance equations for each of the components are solved along streamlines, generated from the pressure field using Darcy's law.

We show that accurate modeling of phase behavior can be achieved if high-order upwind schemes are used along streamlines. The high order of the schemes allows the use of coarser grids along streamlines, thus reducing flash calculations. Also, the longitudinal numerical diffusion is reduced, which is especially important in compositional problems because it can significantly affect the predicted displacement efficiencies. Numerical smoothing is still present as a result of mapping solution values between streamlines and pressure grid at pressure updates. To alleviate this problem, we designed a high-order mapping algorithm that greatly reduces these mapping errors. Besides these recent advances, I will also discuss our plans for the near (and far) future.

Speaker's Web page:
http://ekofisk.stanford.edu/faculty/gerritsen/

Institution Web page:
http://www.stanford.edu/

TOC

December 9, 2003

A Hybrid High-Order Method of the Incompressible Navier-Stokes Equations

Speaker

Jonas Nilsson, Stanford University, jnilson@stanford.edu

A hybrid high-order method of the incompressible Navier-Stokes equations has been developed in an ongoing project between the Department of Scientific Computing at Uppsala University and the Department of Mechanics at Royal Institute of Technology in Stockholm. It is based on fourth-order compact finite difference approximations of Padé type in two dimensions, combined with a spectral method based on a Fourier expansion in the spanwise direction. The Padé operators allow for the use of curvilinear staggered grids. In time, we use a second order semi-implicit scheme, where the nonlinear terms are treated explicitly. In every time step, a system of linear equations is solved for the velocity and the pressure by an outer and an inner iteration. The convergence properties of the iterative method are analyzed. The order of the method is demonstrated in numerical experiments—compute the flow in a channel, the lid-driven cavity, a constricting channel and past a circular cylinder.

Institution Web page:
http://www.stanford.edu

TOC

December 12, 2003

Early Experience in Splintering Communication Protocols

Speaker

Barney Maccabe, University of New Mexico, maccabe@cs.unm.edu

Modern computing systems offer a variety of contexts for execution of code. In splintering, functionality that is traditionally centralized and executed in a single context is broken into small pieces called splinters, which are then distributed among the execution contexts in the system. The goal is to identify a splintering of the functionality so that splinters can be distributed in a way that improves overall system performance while retaining the integrity of the original implementation.

To date, we have focused our efforts on splintering related to communication protocols across the execution contexts provided by a host processor and programmable network interface card in a computational cluster. In this context, splintering is related to OS-bypass and protocol offloading. In the talk, I will describe how splintering differs from these approaches.

Speaker's Web page:
http://www.cs.unm.edu/~maccabe/

Institution Web page:
http://www.unm.edu/

TOC

December 12, 2003

The Stanford Data Stream Management System

Speaker

Jennifer Widom, Stanford University, widom@cs.stanford.edu

This talk will describe our ongoing work developing the Stanford Stream Data Manager (STREAM), a system for executing complex continuous queries over multiple continuous data streams. The STREAM system supports a declarative query language and copes with high data rates and query workloads by providing approximate answers when resources are limited and adapting its execution strategies automatically as conditions change. We will provide an overview of the system and our research plans, highlight several specific contributions to date, and show a demo (time and logistics permitting).

Speaker's Web page:
http://www-db.stanford.edu/~widom/

Research Web page:
http://www-db.stanford.edu/stream/

Institution Web page:
http://www.stanford.edu/

TOC

December 17, 2003

A Data Grid Framework for Managing Planetary Science Data

Speaker

Daniel Crichton, Jet Propulsion Laboratory, crichton@pop.jpl.nasa.gov

As the volume of planetary science data expands seemingly without limit, the ability to share and correlate data across geographically distributed heterogeneous repositories remains a serious challenge. An architectural software framework has been developed by the Object Oriented Data Technology (OODT) task to address this challenge and has been successfully deployed in several diverse scientific domains.

Characterized by separate technology and data architectures, this framework was used by the Planetary Data System (PDS) to enable timely delivery of 2001 Mars Odyssey data to the science community as soon as the data was released from the Mars Odyssey project. As a result, OODT is now the principal infrastructure for delivery of planetary data to the scientists for all future missions and is working with the Planetary Data System to address the huge data volume increases expected for the Mars Reconnaissance Orbiter (MRO).

Based on the success deployment of OODT to support access and management of planetary science data, the Space Physics Archive Search and Extract (SPASE) effort is considering the framework to support a data search and retrieval system for the Space Physics science community. Further demonstrating its flexibility, the framework has been infused into the National Cancer Institute's Early Detection Research Network, enabling access and sharing of data critical to the development of cancer biomarkers across 10 research institutions. In 2003, OODT was selected as Runner Up for NASA Software of the Year.

Research Web page:
http://oodt.jpl.nasa.gov/oodt-site/index.html

Institution Web page:
http://www.jpl.nasa.gov

TOC

December 17, 2003

Merrimac: Supercomputing with Streams

Speaker

William Dally and Patrick Hanrahan, Stanford University,
billd@csl.stanford.edu, hanrahan@cs.stanford.edu

Merrimac uses stream architecture and advanced interconnection networks to give an order of magnitude more performance-per-unit cost than cluster-based scientific computers built from the same technology. Organizing the computation into streams and exploiting the resulting locality using a register hierarchy enables a stream architecture to reduce the memory bandwidth required by representative applications by an order of a magnitude or more. Hence, a processing node with a fixed bandwidth (expensive) can support an order of magnitude more arithmetic units (inexpensive). This in turn allows a given level of performance to be achieved with fewer nodes (a 1-petaflops machine, for example, with just 8,192 nodes) resulting in greater reliability, and simpler system management. We sketched the design of Merrimac, a streaming scientific computer that can be scaled from a $20,000, 2-teraflops workstation to a $20 million, 2-petaflops supercomputer and present the results of some initial application experiments on this architecture.

Speaker's Web pages:
William Dally
http://csl.stanford.edu/~billd/
Patrick Hanrahan
http://www-graphics.stanford.edu/~hanrahan/

Research Web page
http://merrimac.stanford.edu/

Institution Web page:
http://www.stanford.edu/

TOC

From Security to Cells: Ongoing Machine Learning Research at the University of New Mexico

Speaker

Terran Lane, University of New Mexico, terran@cs.unm.edu

In this talk, I will give overviews of a number of ongoing machine-learning projects that my research group is engaged in at the University of New Mexico. I will focus primarily on three projects: Bayesian methods for intrusion detection, causal activation network reconstruction for neuroimaging data, and quantitative modeling of the RNA interference process.

The intrusion detection system (IDS) project stems out of my previous research and attempts to focus on a number of questions that have not been closely examined to date in the adaptive IDS literature. Specifically, we are interested in a semi-supervised Bayesian sensor fusion approach that represents a combination of the traditionally distinct misuse and anomaly detection views of IDS. In the neuroimaging study, we are attempting to reconstruct the neurofunctional underpinnings of cocaine dependence by isolating causal activation networks from functional magnetic resonance imaging (fMRI) data. Similar to genetic regulatory network reconstruction problems, this is a high-dimension Bayesian network structure search problem and is dramatically intractable.

I will briefly review our efforts to incorporate substantial domain knowledge into the process to constrain the search and improve plausibility of the results. Finally, I will describe a new project in which we are attempting to develop a quantitative model of the recently discovered immunological process of RNA interference (a.k.a., RNA silencing, posttranscriptional gene silencing, quelling, etc.) This process is quite exciting because it offers efficient knockdown of specific genes and may represent our first tool for directly attacking viral infections, transposons, and genetic diseases. I will briefly review this process, indicate some of the open questions, and describe potential roles for computer science and machine learning in the analysis of this process.

Speaker's Web page:
http://www.cs.unm.edu/~terran/

Research Web page:
http://www.cs.unm.edu/~terran/research/publist.shtml

Institution Web page:
http://www.unm.edu/

January 9, 2004

Approximate Query Processing with Sampling and Pre-Aggregation

Speaker

Christopher Jermaine, University of Florida, cjermain@cise.ufl.edu

Approximation is an important data management tool, particularly when the data are so numerous that computing an exact answer to a query can be prohibitively expensive. One of the simplest and most powerful approximation techniques available is random sampling. However, sampling can provide very poor approximation quality if the importance of the various data objects is not uniform.

For example, if we are trying to estimate the average net worth of all individuals in the United States, Bill Gates and Warren Buffet cause serious problems for sampling. Standard statistical techniques aimed at overcoming this problem (such as stratification and biased sampling) are often not useful in a database environment where ad-hoc queries must be supported. In this talk, I will describe an interesting database-oriented alternative for boosting the accuracy of sampling over skewed data.

Speaker's Web page:
http://www.cise.ufl.edu/~cjermain/

Institution Web page:
http://www.ufl.edu/

TOC

January 12, 2004

Hybrid Systems: Discontinuous Dynamics in a Continuous World

Speaker

Thomas Seidman, University of Maryland, Baltimore County

We consider the adaptation of ODE theory to situations in which the state space is a hybrid of continuous and discrete components so one switches abruptly from one mode to another as discrete events occur. While the talk primarily considers some general background (and open problems), one fairly concrete example is an optimal control problem involving a bacterial population that switches discontinuously between `dormant' and `active' modes, depending (hysteretically) on the concentration of a critical nutrient. One wishes to have the bacteria consume some pollutant while minimizing the amount of nutrient supplied.

Speaker's Web page:
http://www.math.umbc.edu/~seidman

Institution Web page:
http://www.umbc.edu/

TOC

January 16, 2004

Three Brown Mice: See How They Run – Monitoring Rodent Behavior in the Smart Vivarium

Speaker

Serge Belongie, University of California, San Diego, sjb@cs.ucsd.edu

We address the problem of tracking multiple, identical, nonrigid moving targets through occlusion for purposes of rodent surveillance from a side view. Automated behavior analysis of individual mice promises to improve animal care and data collection in medical research. In our experiments, we consider the case of three brown mice that repeatedly occlude one another and have no stable trackable features. Our proposed algorithm computes and incorporates a hint of the future location of the target into layer-based affine optical flow estimation. The hint is based on the estimated correspondences between mice in different frames derived from a depth-ordering heuristic. Our approach is simple, efficient and does not require a manually constructed mouse template. We demonstrate encouraging results on a challenging test sequence containing multiple instances of severe occlusion. (This is joint work with Kristin Branson and Vincent Rabaud.)

Speaker's Web page:
http://www.cse.ucsd.edu/facresearch/facultyprofiles/BelongieS.html

Institution Web page:
http://w.ww.ucsd.edu/

TOC

January 23, 2004

Type-Based Specialization in a Telescoping Compiler for MATLAB

Speaker

Cheryl McCosh, Rice University, chom@cs.rice.edu

Telescoping languages is a strategy to automatically generate highly optimized, domain-specific libraries. The key idea is to create specialized variants of library procedures through extensive offline processing. Calls to the library in user scripts can then be replaced with calls to the appropriate variants. This talk will describe a telescoping system, called LibGen, that generates high-performance Fortran or C libraries from prototype MATLAB code.

LibGen uses variable types to guide variant generation and specialization. To solve the type-inference problem, LibGen uses a novel strategy that constructs, at each point in the MATLAB routine, a type jump-function describing the types of the local variables in terms of the types of the inputs. The type jump-functions are computed using a propositional, constraint-based formulation of the type-inference problem and an efficient, graph-theoretical algorithm to determine the solution.

This talk demonstrates the power of the approach by showing that LibGen is able to determine all the required variants for the ARPACK library, from MATLAB development code provided by the ARPACK authors. By allowing library writers to develop and maintain their code in higher-level languages, such as MATLAB, while achieving the performance of coding in lower-level languages, such as Fortran or C, telescoping languages strive to increase the productivity of the scientific community.

Institution Web page:
http://www.rice.edu/

TOC

January 26, 2004

On the Farm: Parallel Smalltalk for Simulating Dairy Operations

Speaker

Elizabeth Post, Lincoln University, poste@lincoln.ac.nz

This talk is mainly about porting a complex system with a developed framework and components in several languages from Windows to Linux and at the same time, parallelizing it for a distributed cluster system. The core framework was implemented in Smalltalk whose virtual machine doesn't parallelize easily. This ported and parallelized model was then used for a trial optimization study using a genetic algorithm to identify the best ways of managing dairy farms. Then we had to visualize the large amounts of high-dimensional data from thousands of simulations with eight varying input parameters and one output result.

I will be talking about porting, parallelizing, using it for optimization and visualizing the results. All of this is preliminary development work—the dairy farm model is still being improved and refined to give more realistic results, although it gives some pretty reasonable results already.

Speaker's Web page:
http://www.lincoln.ac.nz/amac/profiles/poste.htm

Institution Web page:
http://www.lincoln.ac.nz/

TOC

January 27, 2004

What's New with SCIRun2?

Speaker

Steven Parker, University of Utah, sparker@cs.utah.edu

I will present an overview of a new parallel component framework (SCIRun2) and will demonstrate how software components are being used in a variety of research projects at the University of Utah. SCIRun2 is based on the Common Component Architecture (CCA) and the Scientific Computing and Imaging Institute's SCIRun and Uintah projects. SCIRun2 supports distributed computing through distributed objects. Parallel components are managed transparently over an MxN method invocation and data redistribution subsystem. A meta-component model based on CCA is used to accommodate multiple component models such as CCA, CORBA and Dataflow.

The SCI Insitute at the University of Utah utilizes component-based environments for biomedical computing, computational combustion and other applications. I will present how components are utilized in the DOE ASCI-funded Center for Simulation of Accidental Fires and Explosions (C-SAFE) to achieve scalability on thousands of processors for multi-physics simulations.

Speaker's Web page:
http://www.cs.utah.edu/~sparker/

Research Web page:
http://www.sci.utah.edu/

Institution Web page:
http://www.utah.edu/

TOC

January 29, 2004

An Optimization Technique for Large-Scale Nonlinear Programming

Speaker

Miguel Argaez, University of Texas at El Paso, mar@math.utep.edu

We are interested in the development of efficient and robust optimization algorithms for solving large-scale nonlinear programming problems. The research plans focus on the use of interior-point methods, a path-following strategy, Krylov subspace methods, and globalization strategies. In particular, the quasicentral path notion, which is a relaxation of the central path that excludes the dual condition, is proposed as a central region to guide the iterates to an optimal solution of the problem.

For making progress towards the central region, a generalized augmented Lagrangian function is introduced as a merit function. Furthermore, a new notion of weighted neighborhoods as a measure of closeness to this central region is presented. A preliminary numerical implementation was conducted in MATLAB using the proposed methodology and is working efficiently with small- to medium-size problems, but our goal is to extend this work for solving large-scale problems.

We propose to decouple the linear system associated with the perturbed Karush-Kuhn-Tucker conditions to obtain a smaller system called a saddle point problem. We investigate efficient linear algebra solvers for obtaining an approximate solution of the reduced system that fits with the proposed interior-point Newton globalization strategy. In collaboration with a National Laboratory, we plan to develop high-quality, portable software of the proposed algorithm incorporating object-oriented software libraries. In particular, one goal is the application of the algorithm to problems of interest to the Department of Energy.

Speaker's Web page:
http://www.math.utep.edu/Faculty/mar

Institution Web page:
http://www.utep.edu/

TOC

January 29, 2004

A Global Optimization Technique for Solving Zero or Very Small Residual Nonlinear Least-Squares Problems

Speaker

Leticia Velazquez, University of Texas at El Paso, leti@math.utep.edu

We consider global minimization of nonlinear least-squares problems with zero or very small residuals. We present an algorithm based on Levenberg-Marquardt methods combined with a multi-start technique. We compare the algorithm numerically with several existing globalization techniques, and our preliminary numerical results demonstrate that the method is quite competitive. The set of problems consists of test cases from the literature and of one well-known model problem in computational biology—the distance geometry problem. We propose to extend this approach for solving constrained nonlinear least-squares problems with nonzero residuals and to apply the algorithm to DOE data-fitting applications.

Speaker's Web page:
http://www.math.utep.edu/Faculty/leti

Institution Web page:
http://www.utep.edu/

TOC

February 10, 2004

Designing for Usability

Speaker

James Hobart, Classic Systems Solutions, Inc., jimh@classicsys.com

This presentation will explore the principles of user-centered design and how to apply these principles to your Web site or application.

Attendees will learn how to apply a proven process for identifying true user requirements, develop and validate conceptual models, and create designs that are highly usable. We will also explore the creation and usage of visual design patterns for solving complex user interface design issues.

Attendees will learn how to:

- Apply principles of user-centered design
- Use proper layout and design techniques
- Use new UI design modeling techniques
- Create and implement in-house Web standards
- Design more successful applications
- Validate and defend important design decisions
- Document design patterns to leverage what your team has learned

Speaker's Web page:
http://classicsys.com/classic_site/html/president.html

Institution Web page:
http://classicsys.com/

TOC

Statistical Condition Estimation

Speaker

Alan Laub, University of California, Davis, laub@ucdavis.edu

Understanding the condition (or sensitivity) of problems solved with algorithms implemented in floating-point arithmetic is an essential step in assessing the accuracy of computed solutions. Standard approaches to measuring the condition of various problems in numerical linear algebra, for example, compress all sensitivity information into a single condition number. Thus, a loss of information can occur in situations where this standard condition number does not accurately reflect the actual sensitivity in a solution or in particular entries of a solution.

A method is described that overcomes these and other common deficiencies. The new procedure measures the effects on the solution of small random changes in the input data, and by properly scaling the results, it obtains condition estimates from each entry of a computed solution. This approach, which is referred to as small-sample statistical condition estimation (SCE), applies to both linear and nonlinear problems. In the former case, when an explicit Frechet derivative is available for the computed quantity in question, the method is especially efficient, costing no more than standard normwise or componentwise estimates. Moreover, SCE has the advantage of considerable flexibility. For example, it easily accommodates restrictions on or structures associated with allowable perturbations (symmetry, bandedness, etc.). The method has a rigorous statistical theory available for the probability of accuracy of the condition estimates. The theory of SCE is described along with several illustrative examples.

Speaker's Web page:
http://www.cs.ucdavis.edu/people/faculty/laub/index.html

Institution Web page:
http://www.ucdavis.edu/

TOC

February 11, 2004

PLAnetary Scale Monitoring Architecture (PLASMA)

Speaker

Demet Aksoy, University of California, Davis, aksoy@cs.ucdavis.edu

Recent advances in wireless sensor networks enable various geo-, air- and water-based monitoring. However, monitoring events, such as disasters or environmental change, requires much more than a local deployment of a sensor network. In order to correlate the observations made by sensor networks, there is an emerging need to exchange mass amounts of globally distributed data.

Our project, PLAnetary Scale Monitoring Architecture (PLASMA), aims at developing an integrated data management and communications system for time-critical response based on real-time observations made by heterogeneous sensors. PLASMA makes use of a multi-tiered, satellite-based architecture to enable the most suitable communication for the application.

In this talk, I will present our energy-efficient and dynamic-channel allocation algorithm, which creates broadcast schedules without any initial knowledge of the network topology. Our results suggest that our online algorithm is competitive to the offline random vertex coloring algorithm that assumes perfect knowledge of the network topology while significantly improving power efficiency.

Speaker's Web page:
http://www.cs.ucdavis.edu/~aksoy/

Research Web page:
http://www.cs.ucdavis.edu/~aksoy/addtnl/publication.html

Institution Web page:
http://www.ucdavis.edu

TOC

February 18, 2004

A Numerical Method for High-Speed Reactive Flow on Overlapping Grids

Speaker

Donald Schwendeman, Rensselaer Polytechnic Institute, schwed@rpi.edu

A numerical method for the accurate simulation of high-speed reactive flow is presented. The method is an unsplit, Godunov-type, shock-capturing scheme for the reactive Euler equations, and uses overlapping grids to handle complex geometries together with a scheme of adaptive mesh refinement (AMR) to locally increase grid resolution. The method has been used to study a number of problems involving high-speed reactive flow focusing on the birth, propagation and failure of detonation waves in a variety of geometries and using a variety of reaction and equation of state models. The talk will provide an overview of the method and will discuss several calculations for a variety of reactive flow problems.

Speaker's Web page:
http://eaton.math.rpi.edu/Faculty/Schwendeman/home.html

Institution Web page:
http://www.rpi.edu/

February 20, 2004

PROLINKS: A Database of Co-Evolving Proteins

Speaker

Matteo Pellegrini, University of California, Los Angeles, matteope@mbi.ucla.edu

Traditionally, protein evolution has been studied by measuring the sequence similarity within groups of homologous proteins. However, it is also possible to measure the co-evolution of pairs of non-homologous proteins. To infer co-evolution, we use several methods:

1. The Phylogenetic Profile technique measures the likelihood that pairs of genes are present or absent together in fully sequenced genomes
2. The Rosetta Stone technique searches for protein-fusion events
3. The gene neighbor technique looks at the intra-genic distance between pairs of genes on multiple genomes.

We will discuss statistical methods to evaluate the likelihood of co-evolution using these three approaches. We will also discuss the observation that co-evolving proteins tend to participate in common biological processes.

Institution Web page:
http://www.ucla.edu/

TOC

February 25, 2004

Adaptive Resource Management via Processor Virtualization: Charm++ and AMPI

Speaker

Laxmikant Kale, University of Illinois at Urbana-Champaign, kale@cs.uiuc.edu

Processor virtualization and consequent message-driven execution has emerged as an extremely powerful technique in our research over the past decade. The basic idea is simple: allow the program to be expressed in terms of a large number of virtual processors and leave their assignment to physical processors on the runtime system (RTS). Yet, this simple idea empowers the runtime system to implement many adaptive strategies that enable automatic checkpointing, flexible reassignment of processors among multiple jobs, out-of-core execution, etc. Most importantly, it allows the RTS to learn the application behavior at runtime and automatically optimize performance via dynamic load balancing and communication optimizations.

These strategies have been embodied in Charm++ and Adaptive MPI (AMPI), which have been used to effectively parallelize several applications, including NAMD, a molecular dynamics program that shared the Gordon Bell prize at SC'02. I will describe Charm++/AMPI and the adaptive runtime system and present an overview of the applications being programmed with it. I will also summarize ancillary research projects based on these ideas, such as domain-specific frameworks, fault tolerance, meta-scheduling on multiple clusters with market-based bartering mechanisms.

Speaker's Web page:
http://www.cs.uiuc.edu/directory/directory.php?name=kale

Research Web page:
http://charm.cs.uiuc.edu/

Institution Web page:
http://www.uiuc.edu/

TOC

March 3, 2004

Speaker

How to Search a Social Network

Lada Adamic, Hewlett Packard Laboratories, ladamic@hpl.hp.com

We address the question of how participants in a small-world experiment are able to find short paths in a social network using only local information about their immediate contacts. We simulate such experiments on a network of actual email contacts within an organization, as well as on a student social networking Web site.

On the email network, we find that small-world search strategies using a contact's position in physical space or in an organizational hierarchy relative to the target can effectively be used to locate most individuals. However, we find that in the online student network, where the data is incomplete and hierarchical structures are not well defined, local search strategies are less effective. We compare our findings to recent theoretical hypotheses about underlying social structure that would enable these simple search strategies to succeed and discuss the implications to social software design.

(This is joint work with Eytan Adar at HP Labs.)

Speaker's Web page:
http://www.hpl.hp.com/personal/Lada_Adamic/

Institution Web page:
http://www.hpl.hp.com/

3D Computation of Laser Cavity Eigenmodes by Finite Elements

Speaker

Christoph Pflaum, Universität Erlangen-Nürnberg, Germany,
pflaum@informatik.uni-erlangen.de

The iteration method of Fox and Li and the Gauss mode analysis are the main method for the numerical simulation of eigenmodes in laser cavities. While the Gauss mode analysis cannot be applied for the computation of complicated 3D eigenmodes, the iteration method of Fox and Li is a real 3D method. Unfortunately, this method often converges very slowly or diverges, e.g., in the case of long resonator geometries. To circumvent this problem, one would like to apply the Finite Element method directly to the Helmholtz equation (or Maxwell equation), which describes the eigenmodes in a laser cavity. The difficulty of this approach is that in the case of long resonator geometries, a large number of finite elements is needed resolve the oscillations of the eigenmodes. Furthermore, the resulting discrete equation system has a very bad condition number.

In view of these problems, we developed a new approach for the numerical approximation of eigenmodes by Finite Elements. One idea of this approach is the ansatz:

$$E(x,y,z) = \exp[i(\bar{k}-\delta)z]u_r(x,y,z) + \exp[-i(\bar{k}-\delta)z]u_l(x,y,z)$$

where $E(x, y, z)$ is the electrical field of the wave and is propagation constant of the free wave, which in the guiding structure is reduced by a small quantity. By this ansatz, the high oscillation of E are factored out, such that u_r and u_l are smooth function and can be approximated by finite elements on a relatively coarse grid. Also, u_r and u_l represent the waves traveling in the right and left direction. Another important advantage of the above ansatz is that the resulting discrete equation system can be preconditioned by relaxations in the direction of the traveling wave. Numerical results and comparisons with the Gauss mode analysis verify the accuracy of our new approach.

Institution Web page:
http://www.uni-erlangen.de/

TOC

March 5, 2004

Middleware Support for Data Ensemble Analysis

Speaker

Joel Saltz, Ohio State University, saltz-1@medctr.osu.edu

Dramatic decreases in the cost of storage, combined with equally dramatic improvements in network connectivity will make it possible for communities to collaboratively generate and analyze very large distributed data sets. We will describe application scenarios in biomedical research, earth science and in climate modeling that motivate this work. These application scenarios will be used to provide a broad view of what advances in systems software are needed to make this vision a reality.

In many application scenarios, data sets describe sensor-acquired or simulated spatio-temporal regions. Our approach is to develop systems software able to leverage spatio-temporal descriptive metadata to support a broad range of application areas. We will describe techniques for optimized distributed data storage, indexing, retrieval and processing of ensembles of spatio-temporal data sets. We will then describe techniques for supporting on-demand data product generation and for handing spatio-temporal and relational queries directed against these data sets distributed among storage systems located in multiple parallel machines and clusters. Finally, we will describe multiple-query optimization techniques that involve grid-based semantic caching along with identification and exploitation of intermediate results shared between the queries.

Speaker's Web page:
http://bmi.osu.edu/personnel/detail.cfm?person_id=27

Research Web page:
http://www.cs.umd.edu/projects/hpsl/chaos

Institution Web page:
http://www.osu.edu/index.php

TOC

March 10, 2004

Actor-Oriented Metaprogramming

Speaker

Stephen Neuendorffer, University of California, Berkeley,
neuendor@eecs.berkeley.edu

The design of embedded software systems is significantly complicated by the system's interaction with the physical world. Concurrent interaction and the passage of time are central to this interaction. Unfortunately, most software system practices today rely on programming languages that primarily represent sequential operations that complete in an unknown amount of time. Attempts to improve the situation through operating system mechanisms, such as interrupts, timers, and threads, often introduce as many new problems as they solve. These mechanisms provide a programmer with significant flexibility, while also exposing many programming pitfalls. One approach to solving this design problem is to provide better programming languages with structured notions of time and concurrency that can be more easily understood.

I describe an actor-oriented system modeling approach that forms the basis for a better way to architect-embedded systems. An actor is a completely encapsulated component that interacts with other actors through explicit communication channels, rather than through shared memory. Because of this restriction, the interaction between actors can be understood entirely through analysis of their interfaces. The interaction between actors can also be orthogonalized from the actors themselves, allowing it to be specified separately through a Model of Computation. We specify embedded systems using Models of Computation that represent time, concurrency, and interaction with the physical world.

Although we model embedded software using behavioral components, it is important to recognize that this does not require a runtime component architecture, such as CORBA. Efficient embedded software can be generated from actor models by specializing generically specified actors to their context in a model. We call this approach to system design "Actor-oriented Metaprogramming" to distinguish it from other component-based techniques. Metaprogramming allows the use of highly generic and reusable actors for design without sacrificing implementation efficiency.

This talk will introduce Actors and Models of Computation and talk specifically about how they are implemented within Ptolemy II. I'll discuss parts of the framework that enable generic actors and give some examples of heterogeneous modeling.

Speaker's Web page:
http://www.eecs.berkeley.edu/~neuendor/

Research Web page:
http://ptolemy.eecs.berkeley.edu/

Institution Web page:
http://www.berkeley.edu/

March 15, 2004

DPOMP: An Infrastructure for Performance Monitoring of OpenMP Applications

Speaker

Luiz DeRose, IBM Research, laderose@us.ibm.com

OpenMP is today's de facto standard for shared memory parallel programming of scientific applications. However, application developers still face a large number of application performance problems, which make it harder to achieve high performance on SMP systems. These problems are difficult to detect without the help of performance tools. Unlike MPI, which has a standard monitoring interface (PMPI), OpenMP does not yet provide a standardized performance monitoring interface. In order to simplify the design and implementation of portable OpenMP performance tools, Mohr et al., proposed POMP, a standard performance-monitoring interface for OpenMP. POMP describes an API to be called by probes inserted into the application by a compiler, a pre-processor, or via a binary or dynamic instrumentation mechanism. With such a performance monitoring interface, users and tools builders can then define their own POMP-compliant libraries for performance measurement of OpenMP applications.

In this talk, I will present a POMP implementation based on dynamic probes. This implementation is built on top of DPCL, an infrastructure for binary and dynamic instrumentation from IBM. The advantage of this approach lies in its ability to modify the binary with performance instrumentation without requiring access to the source code or relinking whenever a new set of instrumentation is required. In addition, I will present two POMP-compliant libraries: Pomprof and the KOJAK POMP library, which provide the functionality for profiling and tracing OpenMP applications respectively.

Institution Web page:
http://www.ibm.com/us/

March 18, 2004

Analyzing Evolution of Virulence through Spatially-Explicit Epidemic Models

Speaker

Boleslaw Szymanski, Rensselaer Polytechnic Institute, szymansk@cs.rpi.edu

Local interactions between individual organisms influence the population dynamics of species and mediate their competition. In this talk, we describe high-performance simulation of virulence in epidemics using spatially explicit, individual-based models of multi-species habitat. We have implemented a cellular automaton model that represents individuals of two competing host species—a parasite serving as a disease vector and a vector borne pathogen. Genetic algorithms are used to model evolution with mutation and crossover operators used to simulate genetic change; we are particularly interested in the evolution of pathogen virulence. Such an approach enables us to track parentage of each organism and to account for heterogeneity of biotic and abiotic (e.g., spatial) aspects of the habitat.

We discuss the implementation of this model on parallel, distributed memory machines (e.g., IBM SP-2, or a computational cluster). The transition executed at each simulation step is complicated by the presence of genetic information in the system state, motivating the design of several parallel algorithms to cope with this model's complexity. One algorithm computes state transition probabilities efficiently based on simultaneous reduction. Another one efficiently implements the mutation process. Yet another, partitions the simulation domain among the processors.

Using the developed system, we tested a hypothesis that competition between pathogen strains drives the evolution of contagious-disease virulence. To this end, we develop a spatially detailed model of coinfection and superinfection. We assume pairwise competition between strains differing in the mortality probability that they impose on infected hosts. Coinfection happens when two or more strains of pathogen infect the same host, while superinfection refers to within-host competition of pathogens in which a more virulent strain can infect a neighboring host already infected by a less virulent strain and displace the less virulent strain from the individual. We let the probability of superinfection depend continuously on the difference in virulence between competing strains. We apply methods of adaptive dynamics to address both convergence stability and evolutionary stability.

The mean-field approximation of the spatial model predicts evolution to criticality: a strain slightly more virulent than the resident strain can always invade and exclude the resident until a further increase in virulence would imply pathogen extinction. Results of a pair approximation to the spatial model depend on the size of the neighborhood over which the pathogen is transmitted. For smaller neighborhoods, the mean-field approximation predicts a continuous evolutionarily stable set; for larger neighborhoods a single virulence is predicted that is evolutionarily and convergently stable. Simulation of the full model suggests that the pair approximation overestimates both the density of infected hosts for given virulence and the stable level of virulence.

The combined analyses indicate that increasing the size of the interaction neighborhood:

1. Increases the maximal virulence that can persist in the absence of strain competition
2. Increases the average virulence experienced by a host population (and so decreases the equilibrium frequency of infected hosts)
3. Increases the range of non-ESS virulence strains capable of ecological coexistence.

Increasing the advantage of superinfection for a given difference in virulence:

4. Increases evolutionarily stable virulence levels
5. However, it reduces the diversity of pathogen strains that may coexist.

Speaker's Web page:
http://www.cs.rpi.edu/~szymansk

Institution Web page:
http://www.cs.rpi.edu/

TOC

March 26, 2004

Topology-Based Exploration of Scalar Fields

Speaker

Gunther Weber, University of California, Davis, ghweber@ucdavis.edu

Trivariate data is commonly visualized using isosurfaces or direct volume rendering. When exploring scalar fields by isosurface extraction, it is often difficult to choose isovalues that convey "useful" information. The significance of visualizations using direct-volume rendering depends on the choice of good transfer functions. Understanding and using isosurface topology can help in identifying "relevant" isovalues for visualization via isosurfaces and can be used to automatically generate transfer functions. Critical isovalues indicate changes in topology of an isosurface—the creation of new surface components, merging of surface components or the formation of holes in a surface component. Interesting isosurface behavior is likely to occur at and around critical isovalues.

Current approaches to detect critical isovalues are usually limited to isolated critical points. Data sets often contain regions of constant value (i.e., mesh edges, mesh faces, or entire mesh cells). We present a method that detects critical points, critical regions and corresponding critical isovalues for a scalar field defined by piecewise trilinear interpolation over a uniform rectilinear grid. We describe how to use the resulting list of critical regions/points and associated values to examine trivariate data.

Speaker's Web page:
http://graphics.cipic.ucdavis.edu/people/profile?pid=67

Institution Web page:
http://www.ucdavis.edu/

April 6, 2004

Performance Analysis and Tuning of MDCASK and PF3D Codes on Itanium Processors

Speaker

Kirk Hays and Max Alt, Intel Corporation

Intel's Itanium processor offers its users the potential to achieve high-instruction and floating-point throughput relative to clock speed through the use of ILP (Instruction Level Parallelism). Making the most efficient use of the Itanium's EPIC (Explicitly Parallel Instruction Computing) architecture, however, can be a nontrivial task for a compiler when code is ambiguous or complex.

In this seminar, two consultants from Intel show how they were able to obtain speedups ranging from 2x to more than 4x the base-optimized case with LLNL's MDCASK and PF3D codes by applying some basic optimizer pragmas and minor code changes. In nearly all cases, actual code changes are generic and can be retained on non-EPIC architectures. The optimizations presented can be directly applied to these codes' contribution to the Science Runs phase on the Thunder cluster.

The speakers will cover the following topics: project summary, runtime environment overview (including how to run and profile the codes), performance tuning process used, optimization steps, performance characteristics summary, issues encountered during optimization, code runs in parallel environment, and recommendations/ suggestions for future optimizations.

Institution Web page:
http://www.intel.com

TOC

May 6, 2004

GridDB: A Data-Centric Overlay for Scientific Grids

Speaker

David Liu, University of California, Berkeley, dtliu@cs.berkeley.edu

We present GridDB, a data-centric overlay for scientific grid data analysis. In contrast to currently deployed process-centric middleware, GridDB manages data entities rather than processes. GridDB provides a suite of services important to data analysis—a declarative interface, type-checking, interactive query processing, and memorization. We discuss several elements of GridDB, including data model, query language, software architecture and query processing, as well as a prototype implementation. We validate GridDB by showing its modeling of real-world physics and astronomy analyses including measurements on our prototype.

Speaker's Web page:
http://www.cs.berkeley.edu/~dtliu/

Research Web page:
http://www.griphyn.org/

Institution Web page:
http://www.berkeley.edu/

TOC

May 14, 2004

Data-Driven Execution of Communication Tolerant Algorithms

Speaker

Scott Baden, University of California, San Diego, baden@cs.ucsd.edu

Many scalable applications organize into distinct phases of communication and computation, and hence, they are amenable to bulk synchronous parallelism (BSP). However, BSP is not a natural model for latency-tolerant applications that employ asynchronous communication to overlap communication with computation. In this talk, I'll present work in progress with a programming model called Tarragon that is intended to facilitate the design of latency-tolerant algorithmic formulations. I'll discuss two applications that motivate the model. The first application is an iterative kernel for finite difference elliptic solvers; the second is a Monte Carlo method for cell microphysiology.

Speaker's Web page:
http://www.cse.ucsd.edu/users/baden/

Institution Web page:
http://www.ucsd.edu/

TOC

May 17, 2004

Globalization Techniques for Newton-Krylov Methods

Speaker

Homer Walker, Worcester Polytechnic Institute, walker@wpi.edu

A Newton-Krylov method is an implementation of Newton's method in which a Krylov subspace method is used to solve approximately the linear subproblems that determine Newton steps. To enhance robustness when good initial approximate solutions are not available, these methods are usually "globalized," i.e., augmented with auxiliary procedures ("globalizations") that improve the likelihood of convergence from a poor starting point. In recent years, globalized Newton-Krylov methods have been used increasingly for the fully coupled solution of large-scale CFD problems.

In this talk, I will review several representative globalizations, discuss their properties, and report on a numerical study aimed at evaluating their relative merits on large-scale 2D and 3D problems involving the steady-state Navier-Stokes equations.

(This is joint work with John Shadid and Roger Pawlowski at Sandia National Laboratories and Joseph Simonis at WPI.)

Speaker's Web page:
http://www.wpi.edu/Academics/Depts/Math/Faculty/walker.html

Institution Web page:
http://www.wpi.edu

TOC

May 24, 2004

Numerical Issues When Modeling Fluid-Elastic Interaction in 3-D with First-Order System Least Squares

Speaker

Jeffrey Heys, University of Colorado at Boulder, heys@colorado.edu

The mechanical interaction between a fluid and solid can be mathematically modeled using a number of different approaches depending on the physical characteristics of the problem being solved. We are interested in systems consisting of a Newtonian fluid, modeled using the Navier–Stokes equations and a linear elastic material with properties similar to a soft tissue. These coupled fluid-elastic problems are inherently nonlinear because the shape of the fluid domain is not known *a priori*, and the computational grid must be moved or mapped.

We typically use elliptic grid generation (EGG) to map the physical domain to a fixed computational domain. A FOSLS formulation of the Navier–Stokes, EGG, and linear elasticity equations provides a number of benefits to solving coupled systems problems, including optimal finite element approximation in a desirable norm (H^1), optimal multilevel solver performance, optimal scalability, and a sharp *a posteriori* error measure.

The optimality and performance of the formulation has been demonstrated extensively in 2D for a variety of problems, including the fully coupled fluid-elastic system. However, as expected, the extension to 3D brings new challenges for both the whole and the individual parts of the coupled system. Some of the issues associated with the extension to 3D have been partially or fully addressed, such as growing complexity in the multilevel solver, iteration schemes between the components of the fully coupled system, extension to parallel computers, and proper scaling of the equations. We are only beginning to answer other questions, including the handling of singularities and p-refinement.

Speaker's Web page:
http://amath.colorado.edu/faculty/heys

Institution Web page:
http://www.colorado.edu/

TOC

May 24, 2004

First-Order System Least-Squares for Problems with Boundary Singularities

Speaker

Chad Westphal, University of Colorado at Boulder, chad.westphal@colorado.edu

Many elliptic boundary value problems have the fortunate property of a guaranteed smooth solution as long as the data and domain are smooth. However, many problems of interest are posed in nonsmooth domains and as a consequence, lose this property at the boundary. In this talk, we consider problems that have nonsmooth solutions at "irregular boundary points," that is, points that are corners of polygonal domains, locations of changing boundary condition type, or both.

Least-squares discretizations, in particular, suffer from a global loss of accuracy due to the reduced smoothness of the solution. We investigate a weighted-norm least-squares method that recovers optimal order accuracy in the weighted functional norm and weighted H^1 norm, and retains L^2 convergence even near the singularity. The method requires only *a priori* knowledge of the power of the singularity, not the actual singular solution. The theory of this general technique is studied in terms of a simplified div-curl system and shown to be similarly effective when applied to other problems. We also investigate the effect of this technique on the algebraic multigrid solvers used on the resulting linear systems.

Speaker's Web page:
http://persweb.wabash.edu/facstaff/westphac/

Institution Web page:
http://www.colorado.edu/

TOC

Induced-Charge Electro-Osmosis

Speaker

Martin Bazant, Massachusetts Institute of Technology, bazant@math.mit.edu

Induced-charge electro-osmosis (ICEO) refers to the nonlinear electrokinetic slip at a polarizable surface when an electric field acts on its own induced double-layer charge. Here, we develop a simple physical picture of ICEO in the context of some new techniques for microfluidic pumping and mixing. ICEO generalizes AC electro-osmosis at micro-electrode arrays to various dielectric and conducting structures in weak DC or AC electric fields. The basic effect produces micro-vortices to enhance mixing in microfluidic devices, while various broken symmetries—controlled potential, irregular shape, non-uniform surface properties, and field gradients—can be exploited to produce streaming flows with small AC voltages. We also present new experiments demonstrating ICEO vortices around platinum posts in polymer microchannels. Such devices may be easily integrated into biomedical microfluidics to reduce mixing times (e.g., for DNA hybridization assays).

Speaker's Web page:
http://math.mit.edu/~bazant

Institution Web page:
http://www.mit.edu/

TOC

June 3, 2004

Structure and Modeling of Weighted Complex Networks

Speaker

Marc Barthelemy, Commissariat a l'Energie Atomique, Marc.Barthelemy@th.u-psud.fr

In addition to topological complexity, real-world networks display a gradation in the intensity strength between nodes—the weights of the links. I will present two examples— the airline connection network and the scientific collaboration network, which are representative of critical infrastructure and social system respectively. These weighted networks exhibit broad distributions and non-trivial correlations of weights that cannot be explained in terms of the underlying topological structure. These results call for the need of the modeling of complex networks, which goes beyond purely topological models. I will present a model that provides an explanation for the features observed in several real-world networks. This model of weighted network formation introduces a dynamical coupling between topology and weights by rearranging weights when a new link is introduced in the system.

Research Web page:
http://arxiv.org/find/cond-mat/1/au:+Barthelemy_M/0/1/0/all/0/1

Institution Web page:
http://www.cea.fr/default_gb.htm

TOC

June 8, 2004

Knowledge Discovery in Networks

Speaker

David Jensen, University of Massachusetts Amherst, jensen@cs.umass.edu

Networks are ubiquitous in computer science and everyday life. We live embedded in social and professional networks, we communicate through telecommunications and computer networks, and we represent information in documents connected by hyperlinks and bibliographic citations. Only recently, however, have researchers developed techniques to analyze and model data about these networks. These techniques build on work in artificial intelligence, statistics, databases, graph theory, and social network analysis, and they are profoundly expanding the phenomena that we can understand and predict.

In recent work, my students and I have discovered a variety of unique characteristics of networks, and we have developed new methods that recognize and exploit these characteristics to produce more accurate and robust statistical models. Emerging applications for these new techniques include citation analysis, Web mining, intelligence analysis, collaborative filtering, computer security, epidemiology, and financial fraud detection.

Speaker's Web page:
http://kdl.cs.umass.edu/people/jensen/

Research Web page:
http://kdl.cs.umass.edu/

Institution Web page:
http://www.umass.edu/

TOC

June 10, 2004

SCons: A Next-Generation Build Tool

Speaker

Steven Knight, SCons Project Founder/Leader, knight@baldmt.com

A next-generation build tool, SCons, simplifies the process of building software reliably and correctly. Implemented in Python, one of the goals of the SCons project is to make software builds easy for non-programmers and programmers alike. The SCons design won the Software Carpentry build tool competition in August 2000, and the project has been attracting attention and users ever since.

The SCons project team emphasizes SCons ability to "do the right thing out of the box" while still offering enough flexibility to solve difficult build problems. This talk will contrast the SCons approach with other build tools like Make and Ant, provide an overview of SCons functionality, and discuss the project's future directions.

Speaker's Web page:
http://www.baldmt.com/~knight/ (broken)

Institution Web page:
http://scons.sourceforge.net/

TOC

Worm Detection and Response: Local Strategies and Analytical Models

Speaker

Wenke Lee, Georgia Institute of Technology, wenke@cc.gatech.edu

Worm detection systems have traditionally focused on global strategies and required a very large monitored network, say 2^{20} nodes. We consider how local networks can provide early detection and complement global monitoring strategies.

In this talk, I will first describe two local-detection approaches. The Destination Source Correlation (DSC) algorithm monitors for worm-like infection and scanning patterns to accurately detect infected hosts. The HoneyStat approach uses modified honeypots to record and analyze system and network events. The analysis result can indicate whether an automated or worm attack is present.

I will also describe our new discrete time-based worm model adapted from the AAWP (Analytical Active Worm Propagation) model. We use this model to evaluate the effectiveness of our local detection algorithms. Our analytical results show that using a /12-monitored network, a worm warning can be issued when only 0.19% of all vulnerable hosts on the Internet are infected.

We also extend our analytical model to evaluate the effectiveness of local response techniques such as traffic throttling or blocking, and patching. Our results show that with 80% deployment ratio of network level local response, Internet worms can be slowed down about five times faster than without local response. If used together with patching, worm propagation can be stopped completely.

Speaker's Web page:
http://www.cc.gatech.edu/fac/Wenke.Lee/

Institution Web page:
http://www.gatech.edu/

TOC

July 21, 2004

Detailed Cache Coherence Characterization for OpenMP Benchmarks

Speaker

Frank Mueller, North Carolina State University, mueller@cs.ncsu.edu

Past work on studying cache coherence in shared-memory symmetric multiprocessors (SMPs) concentrates on studying aggregate events, often from an architecture point of view. However, this approach provides insufficient information about the exact sources of inefficiencies in parallel applications. For SMPs in contemporary clusters, application performance is impacted by the pattern of shared memory usage, and it becomes essential to understand coherence behavior in terms of the application program constructs, such as data structures and source code lines.

We introduce ccSIM, a cache-coherent memory simulator fed by data traces obtained through on-the-fly dynamic binary rewriting of OpenMP benchmarks executing on an SMP node. The novelty of ccSIM lies in its ability to relate coherence traffic—specifically coherence misses as well as their progenitor invalidations—to data structures and to their reference locations in the source program, thereby facilitating the detection of inefficiencies.

Our experiments demonstrate that

a. Cache coherence traffic is simulated accurately for SPMD programming styles as its invalidation traffic closely matches the corresponding hardware performance counters.
b. We derive detailed coherence information indicating the location of invalidations in the application code.
c. We illustrate opportunities for optimizations from these details.

By exploiting these unique features of ccSIM, we were able to identify and locate opportunities for program transformations, including interactions with OpenMP constructs, resulting in both significantly decreased coherence misses and savings of up to 73% in wall-clock execution time for several real-world benchmarks.

Speaker's Web page:
http://moss.csc.ncsu.edu/~mueller/

Institution Web page:
http://www.ncsu.edu/

TOC

On Refractive Optical Flow

Speaker

Sameer Agarwal, University of California, San Diego, sagarwal@cs.ucsd.edu

This paper presents a novel generalization of the optical flow equation to the case of refraction, and it describes a method for recovering the refractive structure of an object from a video sequence acquired as the background behind the refracting object moves. By structure, we mean a representation of how the object warps and attenuates (or amplifies) the light passing through it. We distinguish between the cases when the background motion is known and unknown. We show that when the motion is unknown, the refractive structure can only be estimated up to a six-parameter family of solutions without additional sources of information. Methods for solving for the refractive structure are described in both cases. The performance of the algorithm is demonstrated on real data, and the results of applying the estimated refractive structure to the task of environment matting and compositing are presented.

Speaker's Web page:
http://vision.ucsd.edu/~sagarwal

Research Web page:
http://vision.ucsd.edu/

Institution Web page:
http://www.ucsd.edu/

TOC

July 23, 2004

What Multilevel Parallel Programs Do When You Are Not Watching

Speaker

Gabriele Jost, NASA Ames Research Center, gjost@nas.nasa.gov

The talk will compare three multilevel parallel programming paradigms suitable for shared memory computer architectures—MPI/OpenMP, MLP, and nested OpenMP. The three paradigms are applied to a set of benchmark codes from the field of Computational Fluid Dynamics. A case study will be presented using a performance analysis tool for the comparison. Detailed analysis techniques, made possible by the tool, help to differentiate between the influences of the programming model itself and other factors, such as implementation-specific operating system or architectural issues. Some implementation issues regarding the runtime support for the different programming paradigms will be discussed.

Institution Web page:
http://www.nasa.gov/centers/ames/home/index.html

July 29, 2004

Words and Pictures

Speaker

David Forsyth, University of California, Berkeley, daf@cs.berkeley.edu

It is now possible to obtain huge collections of images that carry annotations of one form or another. These annotations can take many forms. Examples include the keywords that occur in the Corel and Hemera collections using various forms of metadata—who made the artifact depicted, when it was made, etc.—that are common in museum collections; the narrative annotations that are sometimes found in museum collections; and the captions that one can collect with news images.

Such collections are interesting for two reasons. First, because visual and text information tends to be complementary, a relatively simple analysis of both the image and the text can reveal a great deal about the data item. This means that, for example, one can cluster such collections well, enabling naive users to browse a museum's collection or browse the news in a natural way. Second, such collections can be thought of as huge but poorly supervised data sets, containing both information about the appearance of objects and various forms of world knowledge. I will demonstrate a variety of methods whereby one can build probability models linking images or image regions to their annotations. With such models, one can organize a collection in a way that makes browsing easy and quite natural. One can search for pictures using words. And, what is more important, one can attach words to pictures or even to regions. Finally, one can attach names to faces.

Speaker's Web page:
http://www.cs.berkeley.edu/~daf/

Institution Web page:
http://www.berkeley.edu/

TOC

July 29, 2004

ReFlex: Flexible and Reliable Systems Technologies for Responding to Massively Disruptive Events

Speaker

Ling Liu, Georgia Institute of Technology, lingliu@cc.gatech.edu

Emergency response and disaster recovery have always combined two elements: a planned and prepared response to regular disaster events and an ad-hoc, dynamically organized (composed) response to massively disruptive events that exceed the prepared capacity (e.g., great earthquakes, the "Perfect Storm" of 1991, 9/11 attacks). Although these large-scale natural or man-made disasters are rare events, as 9/11 showed, the consequences of such large-scale crises can be far-reaching and diverse. They cause extremely large and lasting damages. This is due partly to difficulties in effective planning, preparation, and execution of responses to mitigate damages of and recover from such enormous disasters.

Economically and politically, it is infeasible to reserve (and keep idle) large amounts of resources and infrastructure in preparation for extraordinary disasters that may or may not happen in our lifetime. Although in the past we have resigned ourselves to a fatalistic apathy towards such rare events, our reaction to the 9/11 attacks show a new resolve to face such adversity.

We propose an innovative approach that combines a low-cost preparation of potentially useful resources and infrastructure components with a dynamically composed flexible and reliable recovery infrastructure. During normal times, resources and infrastructure contribute to economically meaningful purposes. When an extraordinary disaster happens, our technologies can be used to build a flexible, effective, reliable and secure response and recovery infrastructure.

Concretely, we propose to focus on the following three technological aspects for supporting flexible and reliable risk mitigation, emergency response, and disaster recovery. First, we propose a secure overlay network infrastructure and enabling technologies to dynamically construct large-scale, self-configuring, and self-healing overlay networks for emergency response and recovery and to assist the damage containment and recovery efforts. An example of such mechanisms is trust negotiation among multiple emergency response entities. Second, we advocate mechanisms for effective risk mitigation in advance of emergencies by promoting the concept of disaster-resilient communities and techniques for building such communities. Third, but not least, we plan to develop technological solutions to support decision making and optimize the distribution of resources in the presence of inconsistent or maliciously injected conflicting information, aiming at building a flexible, reliable, and robust emergency response and risk recovery network. Flexibility is required for extraordinary disasters since we cannot predict how much damage they may cause. Reliability is required since many enormous disasters have associated secondary risks or crisis (e.g., earthquake aftershocks, correlated terrorist attacks). Robustness is required to provide resilience to coordinated attacks prepared to incapacitate or subvert the emergency response network.

In this talk, I will give a summary of the ReFlex project, a brief overview of the background knowledge on decentralized overlay networks, as well as an in-depth discussion on a selection of technical issues and some of the ongoing efforts of Reflex project at Georgia Tech.

Speaker's Web page:
http://www.cc.gatech.edu/~lingliu/

Research Web page:
http://www.cc.gatech.edu/~lingliu/#ResearchProjects

Institution Web page:
http://www.gatech.edu/

Emergence of Detonation in the Flowfield Induced by Richtmyer–Meshkov Instability

Speaker

Nikolaos Nikiforakis, University of Cambridge, n.nikiforakis@damtp.cam.ac.uk

Combustible mixtures of gases can support two steady modes of combustion, namely, deflagration and detonation. Under certain conditions, a relatively low-speed deflagration can accelerate to form a supersonic detonation wave, a process referred to as deflagration to detonation transition (DDT). Whilst the behaviour of steady deflagrations and detonations is reasonably well understood, there are many gaps in our understanding of the nature of the transition mechanism.

The aim of this research is to investigate the transition process, i.e., the reasons behind the change of propagation mechanism from the advection/reaction/diffusion mode of a deflagration to the coupled shock/reaction system of a detonation wave and in particular the role of interfacial instabilities. To this end, the effect of the Richtmyer–Meshkov instability arising from the interaction of a shockwave with a flame has been studied by means of Implicit Large Eddy Simulations. Transition to detonation is shown to take place in the neighborhood of localised temperature perturbations (hot spots). Finally, the character of the interim combustion-driven waves arising from these hot spots is analysed.

Speaker's Web page:
http://www.damtp.cam.ac.uk/user/nn10005

Institution Web page:
http://www.cam.ac.uk/

TOC

August 13, 2004

Data-Driven Approaches for Biological Networks: Inference, Organization and Analysis

Speaker

Chris Wiggins, Columbia University, chw2@columbia.edu

The emerging disciplines of "systems biology" or "quantitative biology" offer the promise of developing predictive and interpretable models of biological processes that could aid in the design of new biological experiments. The challenge is to extract useful information from large-scale and heterogeneous data sets in the absence of microscopic, few-parameter models.

I will give an overview of current research at Columbia using machine learning and information theoretic approaches for inferring and analyzing networks from large-scale data sets. These include "reverse-engineering" genetic networks by integrating sequence and DNA microarray data, modular organization and visualization of networks, and inference of network growth mechanisms.

Speaker's Web page:
http://www.columbia.edu/~chw2

Institution Web page:
http://www.columbia.edu/

TOC

Efficient Implementation and Parallelization of Mesh-free and Particle Methods — The Parallel Multilevel Partition of Unity Method

Speaker

Alex Schweitzer, Universität Bonn, Germany, m.a.schweitzer@ins.uni-bonn.de

One of the reasons why mesh-free methods gained a lot of attention in recent years in both the engineering and mathematics communities is the fact that mesh-free discretizations and particle models are often better suited to cope with geometric changes of the domain of interest, e.g., free surfaces and large deformations, than classical discretization techniques such as finite differences, finite elements or finite volumes. Another obvious advantage of mesh-free discretizations is their independence from a mesh. Mesh generation is still the most time-consuming part of any mesh-based numerical simulation. However, the cost associated with numerical integration in mesh-free Galerkin methods is usually larger than for mesh-based methods.

There exists a number of variants of mesh-free methods—smoothed-particle hydrodynamics (SPH), reproducing kernel-particle methods (RKPM), element-free Galerkin methods (EFGM), radial-basis functions (RBF), partition of unity methods (PUM), and many more. All these approaches do not depend (at least to a great extent) on a mesh or any fixed relation between the discretization points or particles. The PUM is a very flexible approach due to the product structure of its shape functions. It is a general framework rather than a specific method. The assumptions on the various (independent) components involved in the construction of a PUM function space are abstract properties only, so that we can select problem-dependent components. The PUM approach allows not only for an h-version, p-version and hp-version discretization but also for the exploitation of *a priori* knowledge about the solution in the design of an optimal approximation space.

Some of the key issues involved with mesh-free Galerkin discretization techniques are the fast construction of the shape functions, the numerical integration problem, the treatment of essential boundary conditions, and the efficient solution of the arising linear systems. We will consider these issues in the context of the PUM; however, the presented concepts are applicable to most mesh-free Galerkin approaches.

Institution Web page:
http://www.uni-bonn.de/index_en.shtml

Using Mathematical Models to Understand the AIDS Epidemic and Guide Policy

Speaker

E. Ann Stanley, Los Alamos National Laboratory, spidermoon_nm@yahoo.com

Mathematical models have been used to study the spread of many diseases. These models have provided basic knowledge of epidemic spread and insight into the design of vaccination programs and other control measures. But perhaps the impact of models on policy and research has been greatest in the case of the AIDS epidemic due to its devastating nature and complexity.

I will discuss basic epidemiological theory and then present some insights obtained from AIDS models. These insights have affected study designs for sexual behavior surveys, needle and condom programs, and other aspects of research and public policy. I will finish with a brief discussion of our most recent work on contact tracing, which is one of the most hotly debated methods for controlling HIV spread.

Institution Web page:
http://www.lanl.gov

Mathematical Modeling of Malaria, Early Warning System, and Transgenic Mosquitoes

Speaker

Jia Li, University of Alabama at Huntsville, jli@t7.lanl.gov, li@math.uah.edu

In this talk, I will briefly talk about the life cycle of malaria, present some basic mathematical malaria models, discuss the effects of environmental changes on the transmission dynamics of malaria, the impact of releasing genetically altered mosquitoes in preventing malaria and their preliminary mathematical modeling and analysis.

Speaker's Web page:
http://ultra.math.uah.edu/~li/

Institution Web page:
http://www.uah.edu/

TOC

September 30, 2004

Digital Libraries and Data-Intensive Computing

Speaker

Reagan Moore, San Diego Supercomputer Center, moore@pop.sdsc.edu

Scientific data collections that represent the digital holdings of a research community are now being assembled into digital libraries. Scientists use the digital libraries to support browsing of registered material, discovery of relevant digital entities, and display of the data. This is similar to traditional services provided by digital libraries for image and document collections. However, scientists also need the ability to support manipulation of entire collections as part of data intensive computing. Entire collections are accessed for analysis, streamed through a processing pipeline, and the results are registered back into the digital library. The additional capabilities required by digital libraries to enable data intensive computing are examined for analysis of scientific data collections.

Speaker's Web page:
http://www.sdsc.edu/~moore/rmoore.html

Institution Web page:
http://www.sdsc.edu/

TOC

FlashMob Instant Supercomputing

Speaker

Pat Miller, CASC

It used to be that building a supercomputer of any sort was a multi-million dollar undertaking only available at government laboratories, large corporations, and large academic institutions. Cluster computing in which large numbers of commodity nodes are glued together with a commodity interconnect seems, initially, to have really changed the landscape of parallel computing. However, access to parallel computing is still quite restricted because even owning a small cluster is prohibitive due to several factors, including

1. Cost
2. Experienced personnel to administer it
3. Space
4. HVAC, etc.

What do you do if you want a personal supercomputer? The answer, of course, is a temporary supercomputer in which nodes are volunteered on a temporary basis to form a large-scale, tightly-coupled computing system for the duration of program execution.

We call this new model "FlashMob Computing." A FlashMob supercomputer is constructed from the same building blocks as a conventional cluster—computational nodes and an interconnect. The key differences are that the nodes need not be uniform, as they will be donated by interested parties. The interconnect will be cobbled together from available resources, and there will be limited time to tune the constructed system. Also, in contrast to traditional supercomputing, a FlashMob system is built around a specific application so that non-experts can run significant scientific software on self-built FlashMob clusters.

We'll talk about the "live booting" Linux distro used by FlashMob I and do a live demo of the software. We'll talk about the Supercomputer-in-a-gym event held at the University of San Francisco last April and how the successful packaging of the high-performance LINPACK application has made it possible for relatively unskilled computer enthusiasts to build and benchmark small clusters.

TOC

June 22, 2004

The Time Warp Method of Parallel Discrete Event Simulation

Speaker

David Jefferson, CASC

Time Warp is one of the best known and most general algorithms for scalable parallel discrete event simulation. The simulation model is first decomposed into parallel processes, which then schedule events for one another by sending messages, each timestamped with the simulation time of the event it describes. The fundamental synchronization requirement—the problem that Time Warp solves—is that the simulation as a whole must be executed in such a way that all event messages appear to be processed sequentially, in nondecreasing timestamp order, even though they are neither sent nor arrive in sorted timestamp order, and it is not known statically which processes will send event messages to which others.

Time Warp is a very unusual algorithm in that it relies on a very general form of process rollback for synchronization instead of process blocking. At first glance, general process rollback in a distributed, message-passing environment might seem impossible, or at least extremely difficult, to implement efficiently. However, Time Warp solves the problem in an extremely elegant way using the notion of antimessages—if two antimessages are enqueued in the same queue, they both disappear and the queue gets shorter.

Time Warp was originally implemented at the Jet Propulsion Lab in the early 90s on the Mark III Hypercube and the BBN Butterfly. Performance results on benchmark problems running on up to 112 nodes will be presented. Even now, variations of Time Warp are the best methods known for general parallel discrete simulations without special structure.

TOC

June 29, 2004

BlueGene/L: The Next-Generation of Scalable Supercomputer

Speaker

Kim Yates, CASC

We are giving an overview of the BlueGene/L Supercomputer. This is a jointly funded research partnership between IBM and the Lawrence Livermore National Laboratory as part of the ASC Advanced Architecture Research Program. This massively parallel system of 65,536 nodes is based on a new architecture that exploits system-on-a-chip technology to deliver target peak processing power of 360 teraFLOPS (trillion floating-point operations per second). The machine is scheduled to be operational in early 2005, at price/performance and power consumption/performance targets unobtainable with conventional architectures.

TOC

July 13, 2004

Scientific Data Mining: The Sapphire Project

Speaker

Chandrika Kamath, CASC

The Sapphire project is developing scalable algorithms and software for the interactive exploration of large, complex, multi-dimensional scientific data. We are using ideas from data mining to improve the way in which scientists extract useful information from data. Our work focuses on research in algorithms, incorporation of this research into software, and the application of the software to real-world problems at LLNL. The needs of these applications drive our research. In this talk, I will describe what is involved in data mining.

I will discuss our work using examples from applications, such as detection of human settlements in satellite imagery, detecting and tracking moving objects in video, and finding similar objects in data.

July 20, 2004

Babel Language Interoperability Tool

Speaker

Tom Epperly, CASC

Babel is a state-of-the-art language interoperability tool for high-performance computing being developed at Lawrence Livermore National Laboratory. Babel provides efficient, bidirectional communication between programs in C, C++, Fortran 77, Fortran 90, Java and Python. Compared with industrial alternatives, such as CORBA and COM, Babel's key distinguishing features are high performance, support for Fortran, and support for scientific data types, such as complex numbers and large multi-dimensional arrays.

Babel provides a uniform object model and exception handling across all its supported languages. It uses its own Scientific Interface Definition Language to define inter-language programming interfaces. Internally, Babel achieves interoperability through the use of the intermediate object representation (IOR) in C.

By having a common intermediate representation, Babel converts the many-to-many language interoperability problem into two many-to-one problems. Object-oriented polymorphism is achieved using virtual function table implemented in the IOR. Babel is a key architectural element of the Common Component Architecture—a component model for high-performance computing.

TOC

July 27, 2004

Why Software Quality Assurance Practices Become Evil!

Speaker

Gregory M. Pope, CADSE

This paper looks at the challenge of determining the best practices for software development and why the topic usually sparks a lively debate. The premise is that best practices are application specific and are not easily portable from one industry to another. The paper considers a case study of three different types of software developments and contrasts their differences. The paper presents an alternative method to best practices, which is a common set of principles that are turned into appropriate best practices based on project risk.

TOC

NIF Control System

Speaker

Kim Minuzzo, NIF

The National Ignition Facility (NIF) at LLNL supports the U.S. Department of Energy's National Nuclear Security Administration (NNSA) Defense Programs and LLNL missions of ensuring that the nation's nuclear weapons remain safe, secure, and reliable without nuclear testing. NIF is a 192-beam laser for creating conditions of extreme temperatures and pressures in the laboratory, and when completed, it will be the largest laser in the world.

At LLNL, we are developing the distributed control system required to operate NIF, called the Integrated Computer Control System (ICCS). ICCS uses a distributed object-oriented architecture based on CORBA to provide the interface between the software components distributed across the embedded processors and the supervisory workstations used by the operators to control NIF. User interfaces are coded in Java and most of the control software is written in Ada. The VxWorks real-time operating system is used on the embedded processors. The speaker will present an overview of NIF, ICCS and the SQA practices used on the ICCS.

TOC

August 10, 2004

Multiresolution Computation and Presentation of Topological Structures

Speaker

Valerio Pascucci, CASC

Scalar fields are used to represent data in a wide variety of applications like scientific computing, geographic information systems or medical imaging. Some fundamental topological features of a scalar field are represented in a simple graph structure that is know in different fields as a Contour Tree or Reeb graph.

In this talk, I will present the first technique that allows efficient computing of an augmented version of the Contour Tree and encompasses a complete topological characterization of the 3D scalar field. The new information provided allows the user to know, for example, the number of tunnels and voids of an isosurface without rendering or even computing the isosurface itself. The complexity analysis shows how the new approach allows the computing of the added information with minimal penalty.

The most recent extensions of this work include the decomposition of the Contour Tree in a multi-resolution data structure and its progressive presentation in real time.

I will present specific examples highlighting how an interface based on the progressive contour tree allows the user to understand the "structure" of a scalar field more accurately and minimizes the time spent exploring the data interactively.

TOC

August 17, 2004

Solving Problems with Evolutionary Algorithms

Speaker

Erick Cantú-Paz, CASC

Evolutionary algorithms (EAs) are randomized search methods inspired by evolution and genetics. EAs have been used successfully in many applications for science, engineering, and business. In this talk, I will introduce the basic evolutionary algorithms and present several applications of EAs to data-analysis problems at LLNL. I will also discuss other interesting applications of EAs in art and engineering design.

TOC

June 9, 2004

A Science-Based Case for Large-Scale Simulation

Speaker

David E. Keyes, Columbia University and ISCR

The July 2003 report "A Science-based Case for Large-scale Simulation" (SCaLeS) from the U.S. Department of Energy (DOE) documents new levels of importance for supercomputing in scientific discovery. Together with four other federal reports on supercomputing within the past year, it made the case for an ultrascale computational facility as the DOE's second-highest priority facilities project over the next 20 years, following only the International Thermonuclear Experimental Reactor (ITER, www.iter.org) in importance.

The SCaLeS report takes a look at DOE's Scientific Discovery through Advanced Computing (SciDAC) program—a collection of 51 interconnected projects throughout science, mathematics, and computer science—halfway through its five-year scope, evaluates its early success in mixing the "cultures" of science applications (e.g., astrophysics) with enabling technologies (e.g., adaptive meshing tools), extrapolates the potential of simulation in the eyes of the applications scientists, and outlines hurdles to realizing that potential. It concludes with eight major recommendations. Approximately 315 computational scientists, applied mathematicians, and computer scientists contributed to SCaLeS.

In this talk, the editor of the SCaLeS report summarizes its findings and recommendations, and highlights some of the computational science and software infrastructure of the SciDAC initiative. This software targets multi-teraflop/s platforms and much of it is freely available. Scalable solution algorithms for simulations based on systems of partial-differential equations are emphasized, and are one of the recognized challenges for overall application scalability and of special interest to the speaker.

SCaLeS Web site:
http://www.pnl.gov/scales

Project Web site:
http://www.tops-scidac.org

Speaker Web site:
http://www.columbia.edu/~kd2112/

TOC

Power Presentations

Speaker

Gary Kumfert, CASC

For this talk, think of me as Dear Abbey. Good presentation skills are like good manners. It's something you have to be taught. It's something you have to practice consistently. It's not for your benefit, but for the benefit of people you interact with. It's most effective when it appears natural and effortless. And it says something about either how well educated you are and/or how highly you regard your audience.

Professional researchers are professional presenters. Presentations are a major tool in building your reputation, acquiring funding and defending your work when funding is up for review. This talk will cover slides, props (e.g., podium, pointers), showmanship, and lecture hall etiquette. It's a good refresher even for experienced speakers, and it's good practice for me since students are quick to point out when I don't follow my own advice!

June 30, 2004

Overview of Research Activities in the Center for Applied Scientific Computing

Speaker

Peter G. Eltgroth, Director, CASC

An overview of the Center for Applied Scientific Computing (CASC) is presented. It includes a brief description of CASC activities in advanced software technology, applied mathematics, computational physics, computer science, data science, numerical methods, scalable algorithms for parallel processing, and scientific computing. Some technical detail is provided for ongoing projects in multigrid methods for linear system solution, code transformations, graph theory, adaptive mesh refinement, and computational biology.

TOC

July 6, 2004

Effective Architectures for Scientific Programming

Speaker

Paul Dubois, CASC

Scientific programming is dominated by change. The challenge is to develop scientific simulations in a way that allows them to remain correct despite a high and persistent rate of change not faced by programmers in other disciplines. This talk discusses a variety of ideas related to this challenge.

July 7, 2004

Computational Science at Lawrence Livermore National Laboratory

Speaker

Steven F. Ashby, Director, Computing Applications and Research

Large-scale simulation plays an increasingly important role in many scientific and engineering applications, including those in the defense, energy, and life sciences. The time and length scales of interest in many physical and biological processes span several orders of magnitude, thus requiring the use of sophisticated numerical methods and massively parallel computers.

Dr. Ashby will discuss some of the challenges inherent in terascale simulation with an emphasis on the computational science activities at Lawrence Livermore National Laboratory (LLNL), which currently houses more scientific computing horsepower than any other facility in the world. After giving a brief overview of the Laboratory, he will describe the LLNL simulation environment and highlight some of its computer science and mathematics research efforts. The majority of the presentation will be spent surveying a number of examples of computational science at the leading edge—showing how large-scale scientific simulation is being used to advance scientific discovery. In particular, this presentation will emphasize the importance of combining simulation, theory, and experimentation via illustrative case studies drawn from a variety of scientific and engineering application areas.

TOC

July 14, 2004

Computational Microfluidics

Speaker

David Trebotich, CASC

Microfluidics is the next-generation technology for miniaturization, portability, and networking of the macroscale fluidic systems currently used in "chem/bio" applications. The underlying physics—complex biological flow at the microscale—is not well understood. Also, design of microdevices is lengthy and costly due to imperfections in the microfabrication cycle. The on-demand ability to predict the behavior of biological macromolecules in microenvironments will give significant advantage over current trial-and-error techniques and available rudimentary design tools. The computational microfluidics effort at LLNL will help improve design times and costs, as well as provide an optimization and advanced design platform by developing new algorithms to model complex fluids at the microscale and integrating those algorithms into a design tool.

TOC

July 21, 2004

Radiation Transport in 3D Random Media; Direct Numerical Simulation

Speaker

Frank R. Graziani, B Division

The transport of radiation through materials characterized by large-scale heterogeneities is important in a wide range of applications. In high-energy density physics, it is the propagation of photons through turbulent media. In atmospheric physics, it is solar radiation coupling to the ocean and land mass through clouds. In astrophysics, it is the transport of photons from young stars through turbulent molecular clouds.

In this talk, we will briefly discuss the tremendous theoretical effort that has gone into characterizing this problem. We will show that despite this effort, many problems of relevance to applications are still unsolved. Using the computer as a laboratory, we show how direct numerical simulations of transport through random media allow one to understand and characterize the transport behavior. In particular, using optical-path distribution functions calculated numerically allows one to compute effective mean-free paths and their fluctuations. This tool allows insight into theoretical models and offers hope in understanding the transport of photons in random media in the laboratory.

TOC

July 28, 2004

Architectural Overview of BlueGene/L Supercomputer

Speaker

Don Dossa, CASC

An overview of the hardware and software architecture of LLNL's next supercomputer will be discussed. This 360-teraflop system features several new approaches to massively parallel computing. In addition to describing the system, several areas of applications will also be discussed to show how this system will increase the level of sophistication and breadth of simulations at LLNL.

TOC

August 11, 2004

DOE's Effort to Reduce Truck Aerodynamic Drag — Joint Experiments and Computations Lead to Smart Design

Speaker

Rose McCallen and Kambiz Salari, CASC

At 70mph, overcoming aerodynamic drag represents about 65% of the total energy expenditure for a typical heavy truck vehicle. The goal of our U.S. Department of Energy supported consortium is to establish a clear understanding of the drag-producing flow phenomena. This is being accomplished through joint experiments and computations, leading to the "smart" design of drag-reducing devices. This presentation will describe our objective and approach, provide an overview of our efforts and accomplishments, and discuss our direction for the future.

TOC

August 18, 2004

Developing Interoperable Meshing and Discretization Components

Speaker

Lori Freitag-Diachin, CASC

Typically, the first step in numerically solving a PDE-based application is to generate a discrete representation of the computational domain (the mesh) and approximate the continuous differential operators and solution field on that mesh (the discretization). Over the years, a variety of mesh generation tools and discretization methods have been developed that offer different advantages and disadvantages for different application regimes. In addition, advanced techniques and software tools that provide adaptive mesh refinement, time-varying meshes, mesh-to-mesh data transfer, and parallel decomposition of the mesh have been shown to have significant impact on the application areas that employ them. In each case, the software tools providing these advanced capabilities are becoming increasingly accepted by the scientific community, but it is often not clear *a priori* which techniques are best suited to solve a particular application problem. Ideally, the application scientist should be able to easily insert and experiment with a number of different meshing and discretization software tools, but the application programming interfaces are rarely compatible, making experimentation a labor-intensive and error-prone code modification process.

The Terascale Simulation Tools and Technologies (TSTT) Center has been funded by the Department of Energy's SciDAC initiative to address the barriers preventing easy interoperability and interchangeability of multiple mesh and discretization strategies within a single simulation. We are focusing our effort on the creation of common interfaces for existing TSTT Center technologies that will provide better interoperability and fundamentally increase capabilities that allow application scientists to easily switch among them. I describe the current status of our interface definition effort, the tradeoffs required to balance performance and flexibility, the tools used to address language interoperability issues, and our approach to simplifying the adoption process. To ensure the relevance of our research and software developments, we collaborate closely with both SciDAC application researchers and other technology centers. In particular, I will describe the use of the TSTT interfaces and philosophy to insert advanced adaptive mesh refinement (AMR) and error-estimation procedures into an accelerator-modeling code. This code will develop a new capability that combines front tracking and AMR to deploy a TSTT-compliant mesh quality improvement toolkit into three mesh generation codes in as many days.

TOC

August 25, 2004

Scalable Linear Solvers: Multigrid Methods

Speaker

Rob Falgout, CASC

In this talk, we discuss multigrid methods development at LLNL. Multigrid methods are scalable primarily because they have optimal work requirements. That is, a linear system with N unknowns can be solved by a multigrid method with O(N) work. This property makes it possible to solve ever larger problems on proportionally larger parallel machines in constant time. Classical iterative methods like conjugate gradients are not scalable.

We discuss geometric multigrid methods for solving structure-grid problems, algebraic multigrid methods (AMG) for unstructured-grid problems and parallelization issues for both. The majority of the talk will focus on AMG, a multigrid approach that already works remarkably well over a wide variety of applications. In some cases, however, AMG is not effective without making certain problem-specific modifications and carefully tuning parameters. To address this, CASC researchers have developed new AMG algorithms, new parallelization techniques, and new theory to guide future algorithm development.

June 10, 2004

Data Mining Methods for Network Intrusion Detection

Speaker

Terry Brugger, NAIC

Network intrusion detection systems have become a standard component in security infrastructures. Unfortunately, current systems are poor at detecting novel attacks without an unacceptable level of false alarms. We propose that the solution to this problem is the application of an ensemble of data mining techniques that can be applied to network connection data in an offline environment, augmenting existing real-time sensors.

This talk will expand on this motivation, particularly with regard to operating in an offline environment and our interest in multisensor and multimethod correlation. We'll then review existing systems, from commercial systems to research-based intrusion detection systems. Next, we'll survey state-of-the-art methods in the area. Standard data sets and feature extraction turned out to be more important than we had initially anticipated, so each will be covered in some depth. Next, we will review the actual data mining methods that have been proposed or implemented. We'll conclude by summarizing the open problems in this area, along with some questions of a broader scope. We hope that by providing the motivation and summarizing the work in this area we can stimulate further research.

Terry Brugger is a Ph.D. student at the University of California, Davis, where he's researching Network Intrusion Detection Systems in addition to his day job as a computer scientist for Lawrence Livermore National Laboratory's Information, Operations and Assurance Center. Terry is the principal investigator for LLNL's College Cyber Defenders (CCD) program, which gives college students real-life experience in computer security technologies.

July 1, 2004

Internet Ballistics: Identifying Internet Adversaries Despite Falsified Source Addressing

Speaker

Tony Bartoletti, NAIC

Network scanners constantly probe the Livermore network looking for vulnerabilities. Analyzing packet arrival timing data reveals highly distinctive patterns that may correlate with the attacker's choice of tools, physical platform and/or network location. Consistent identification will improve network security and aid counterintelligence efforts. We have developed tools to preprocess scan data using wavelet techniques to achieve a more than 1,000x compression ratio while still preserving essential features. Initial experiments indicate our methods consistently identify patterns in the data.

TOC

July 8, 2004

A Tale of Two Rootkits — SucKit and Hacker Defender

Speaker

Bill Orvis, EE-EETD

Two different rootkits have been found in compromised systems during the last year; SucKit (SK) and Hacker Defender (HkDef). SK is a Linux rootkit and HkDef is a Windows rootkit. Both of these rootkits implement new and interesting methods of hiding themselves and the intruder's files and processes. They also implement hidden backdoors that would not be detected by normal scanning. This paper describes the operation of the two rootkits, how they hide files and processes, and how they implement a hidden backdoor. It also discusses how to detect these rootkits and how to scan for the hidden backdoor.

TOC